Photo-realistic Renderings for Machines

20105034 Seong-heum Kim

CS580 Student Presentations

2016.04.28

KAIST

Photo-realistic Renderings for Machines

Scene radiances

↓

Light Transport

↓

Pixel values

+ Simulated annotation database.

↓

Machine Learning

Model descriptions
(Light, Shape, Material, Camera)

Physically-based rendering
(Rendering equation with Monte Carlo Integration)

Part 1. Analysis on stochastic sampling errors

Artificial Intelligence
(Convolutional Neural Network)

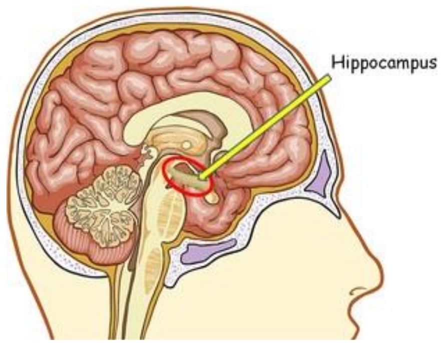

Self-training?
AlphaGo vs AlphaGo

Part 2. One application of this approach

Presentation Papers

Part 1. Analysis on stochastic sampling errors

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration (SIGGRAPH 2013)
Kartic Subr, Jan Kautz

- To study useful indicators for evaluating sampling patterns
- To analyze Gaussian jittered sampling

Part 2. One application of this approach

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views (ICCV 2015)
Hao Su*, Charles R. Qi*, Yangyan Li, Leonidas J. Guibas

- To find a good application of utilizing CG renderings
- To use PBRT results for learning camera viewpoints

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration

Shiny ball in motion

$$\text{Pixel value} = \iiiint \text{Rendering equation } f(x)$$

multi-dim. integral → "sampled" integrand. But,

High variance

High bias

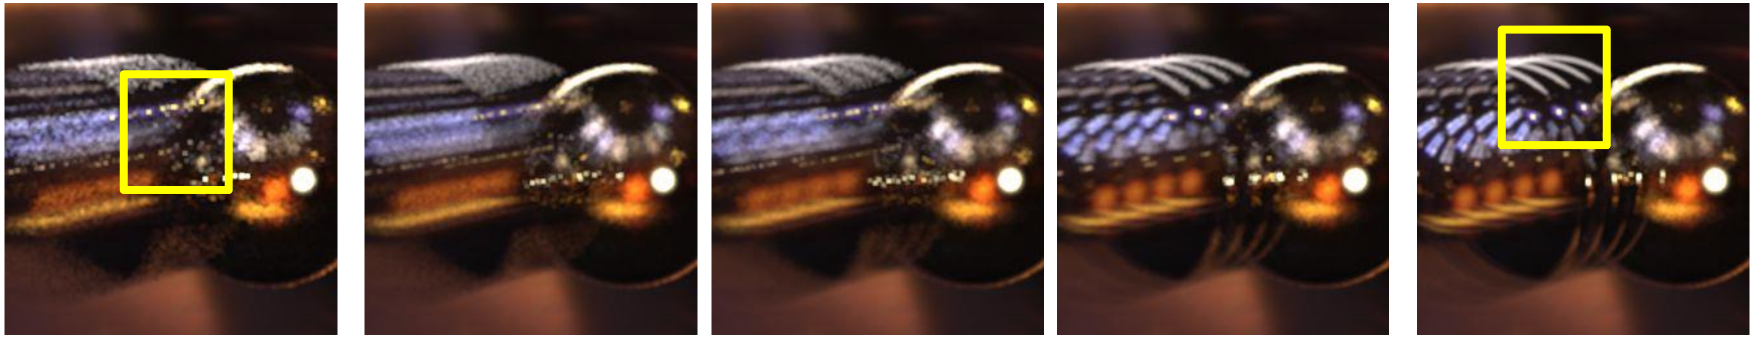

Analysis is non-trivial!

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

Overview of this Paper

- **Stochastic sampling strategies involves random variables.**
  - Accuracy and precision of my estimations?

- **Spectral analysis reduces aliasing effects in estimators**
  - Direct insight for predicting into the first and second order statistics of the integrators. (theoretical contribution)

- **Let's apply it to analyse simple variants of jittered sampling.**
  - Experimental results to show trade-off relationships between random and regular sampling patterns.
  - Quantitative, qualitative comparisons with other strategies (PBRT).

Monte Carlo Estimator

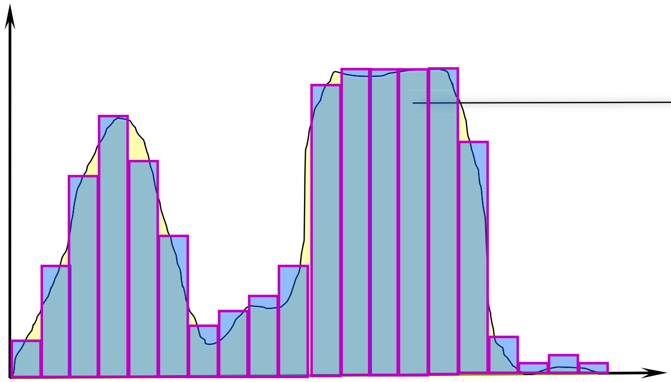

$$I = \int_0^T f(x)dx \approx \frac{1}{N}\sum_{i=1}^{N} \alpha_i f(\boldsymbol{X_i})$$

α_i: normalized weights

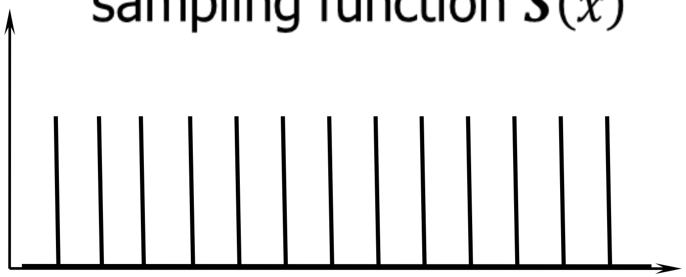

$$= \int_0^T f(x)\boldsymbol{S}(x)dx$$

$$\boldsymbol{S}(x) = \frac{1}{NT}\sum \alpha_i\, \delta(x - \boldsymbol{X}_i)$$

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

KAIST

Strategies to Improve Estimators

1. modify weights

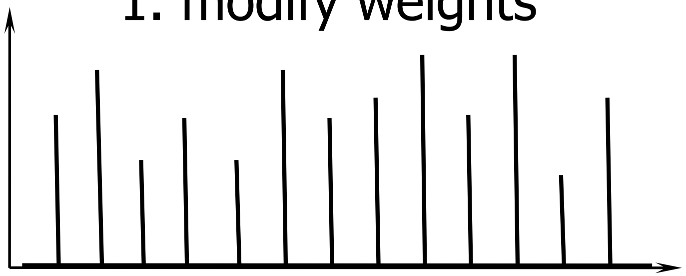

eg. quadrature rules

2. modify locations

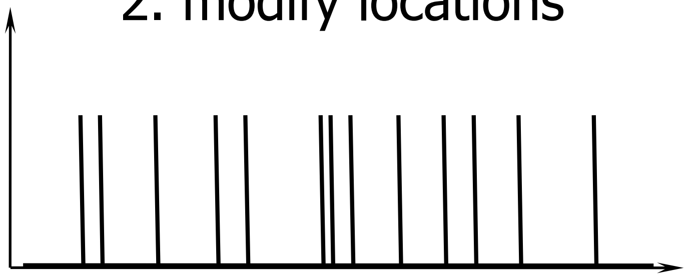

eg. importance sampling

$\alpha_i \delta(x - \boldsymbol{X}_i)$ where $\mathbf{X}_i \sim g(x)$ PDF!!

*** Key idea**

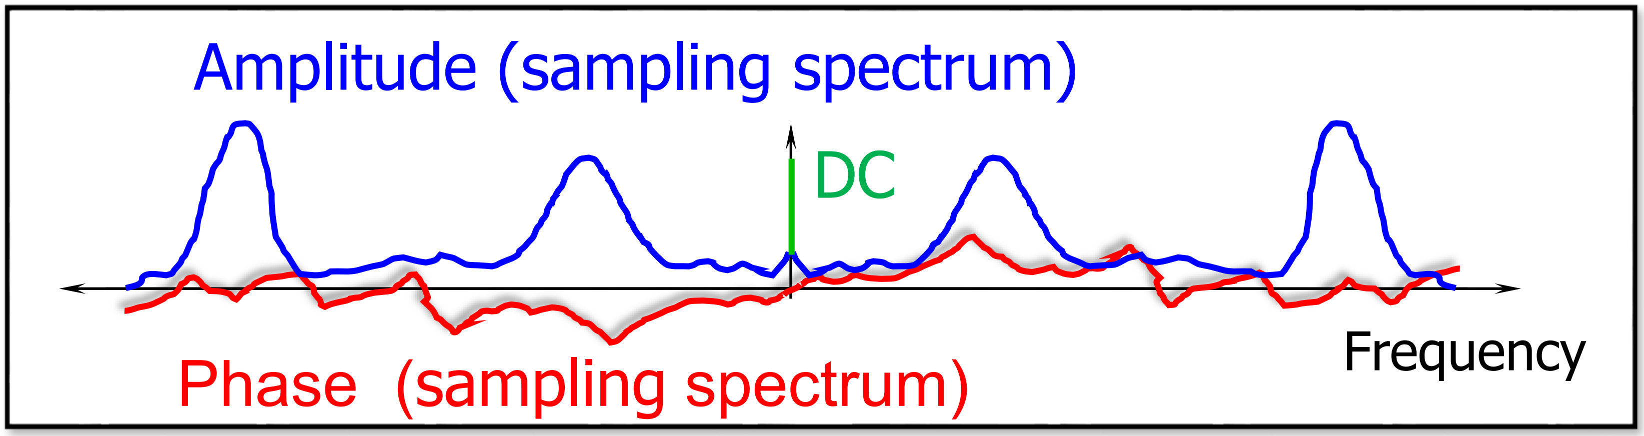

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

Assessing Estimators using Sampling Spectrum

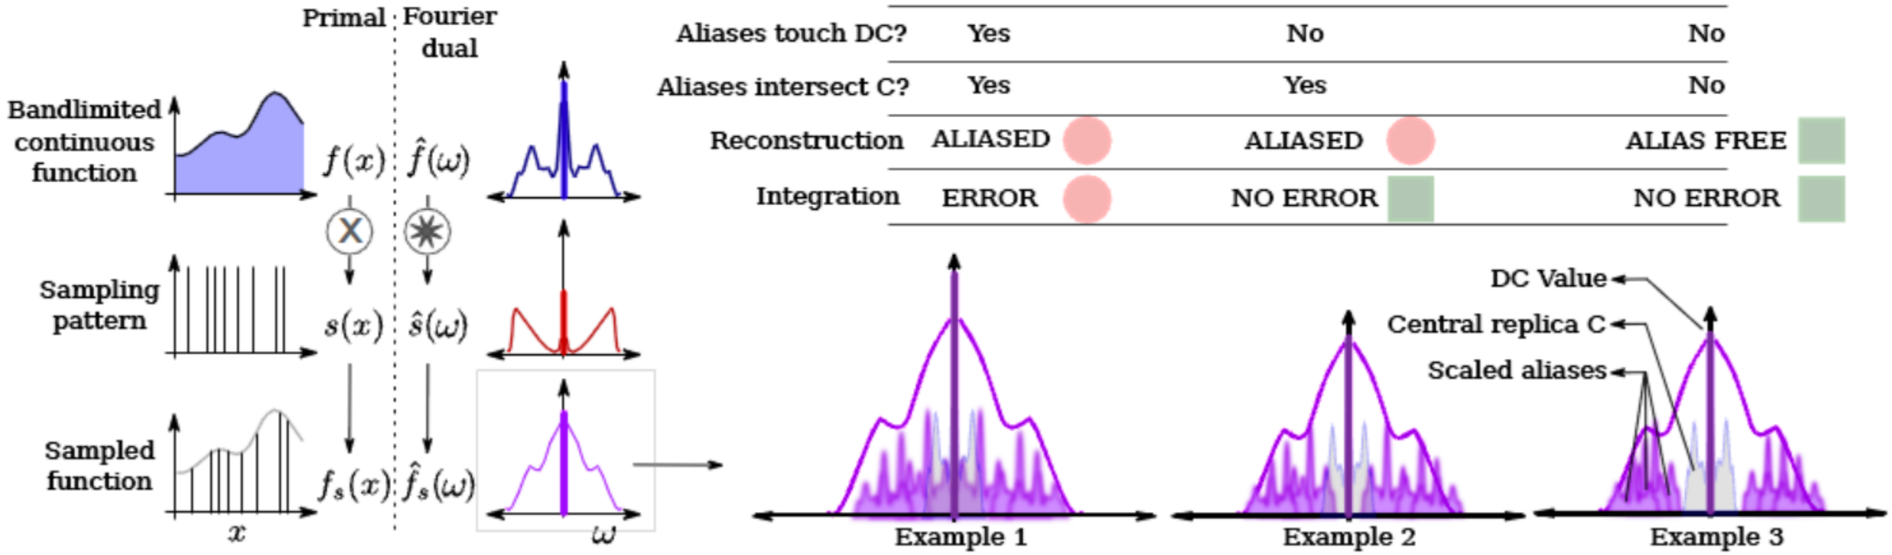

	Example 1	Example 2	Example 3
Aliases touch DC?	Yes	No	No
Aliases intersect C?	Yes	Yes	No
Reconstruction	ALIASED	ALIASED	ALIAS FREE
Integration	ERROR	NO ERROR	NO ERROR

High-level Messages from Fourier domain

Example 1. Low frequency (gain) in sampling spectrum pollutes DC value $\widehat{f_s}(0)$.

Example 2. In high freq., aliased copies need to get more samples (Nyquist rate).

Example 3. A shift (phase) in sampling pattern also changes the aliasing effects.

Bias and variance of secondary MC estimators

Error definition: $\Delta \equiv I - \int_0^T f(x)\boldsymbol{S}(x)dx$

Bias:
(Expected error) $E[\Delta] = <\Delta> = \hat{f}(0) - \int <\boldsymbol{S}(w)> f(w) dw$

Variance: $V(\Delta) = \int V\left(\boldsymbol{S}(w)\right)\left(f(w)\right)^2 dw$

Band-limited
$f(x), S(x)$
↓
$f(w)$
$\boldsymbol{S}(w)$

*Point: Ideal sampling spectrum has no energy in sampling spectrum at frequencies where integrand has high energy

$E[\Delta] \approx E[\boldsymbol{S}(w)]\,(f(w))$ $\quad V(\Delta) \approx V(\boldsymbol{S}(w))\,(f(w))^2$

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

KAIST

Bias and variance of secondary MC estimators

- **Unbiased estimator:** $\langle \hat{\mathbf{S}}(\omega) \rangle = \delta(\omega)$

- **General sampling function:** $\mathbf{S}(x) = \alpha(\mathbf{X}_i)\, \delta(x - \mathbf{X}_i)$

  with $\mathbf{X}_i \sim g(x)$, where $g(x) : [0, T] \mapsto \mathbb{R}^+$ (PDF)

- **Expected Fourier spectrum**

$$\begin{aligned} \langle \hat{\mathbf{S}}(\omega) \rangle &= \langle \alpha(\mathbf{X}_i)\, (\cos(2\pi\omega\mathbf{X}_i) + \imath \sin(2\pi\omega\mathbf{X}_i)) \rangle \\ &= \int \alpha(x)\, (\cos(2\pi\omega x) + \imath \sin(2\pi\omega x))\, g(x)\, \mathrm{d}x \end{aligned}$$

- **Weighting scheme (unbiased):** $\alpha(x) = 1/g(x)$ → PDF
- **Random sampling:** g(x) → constant PDF

Bias and variance of secondary MC estimators

Unbiased sampling,

Imaginary

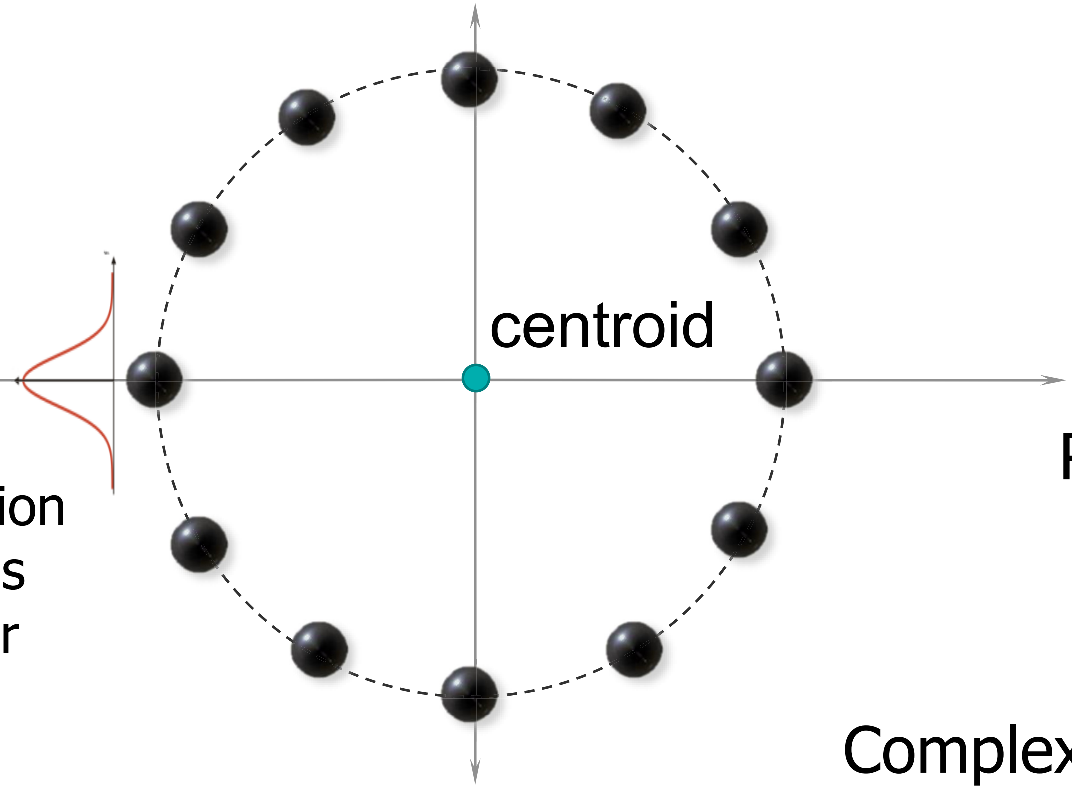

Real

Key idea:
Jitter in position
manifests as
phase jitter

Complex plane

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

KAIST

Case study: Gaussian jittered sampling

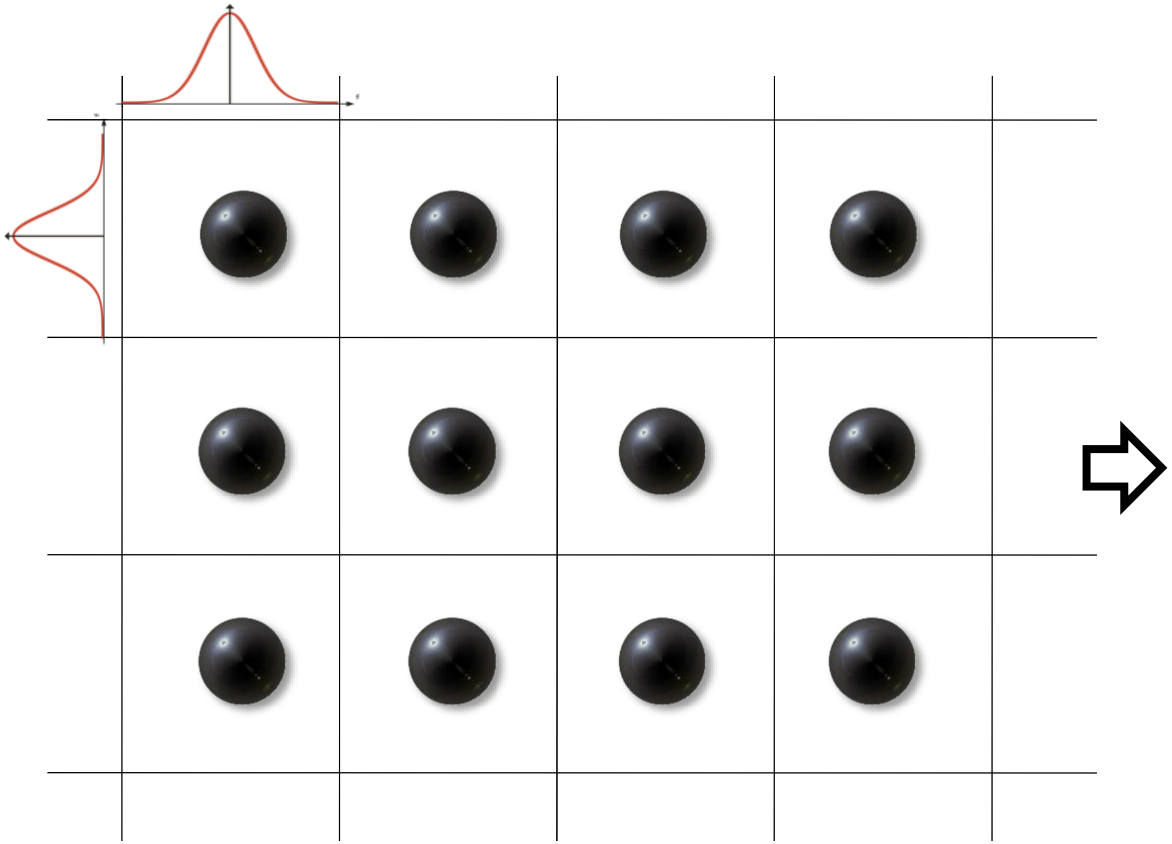

Overall, jitters affect expected centroid.

(More derivations in Appendix.)

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

KAIST

Case study: Gaussian jittered sampling

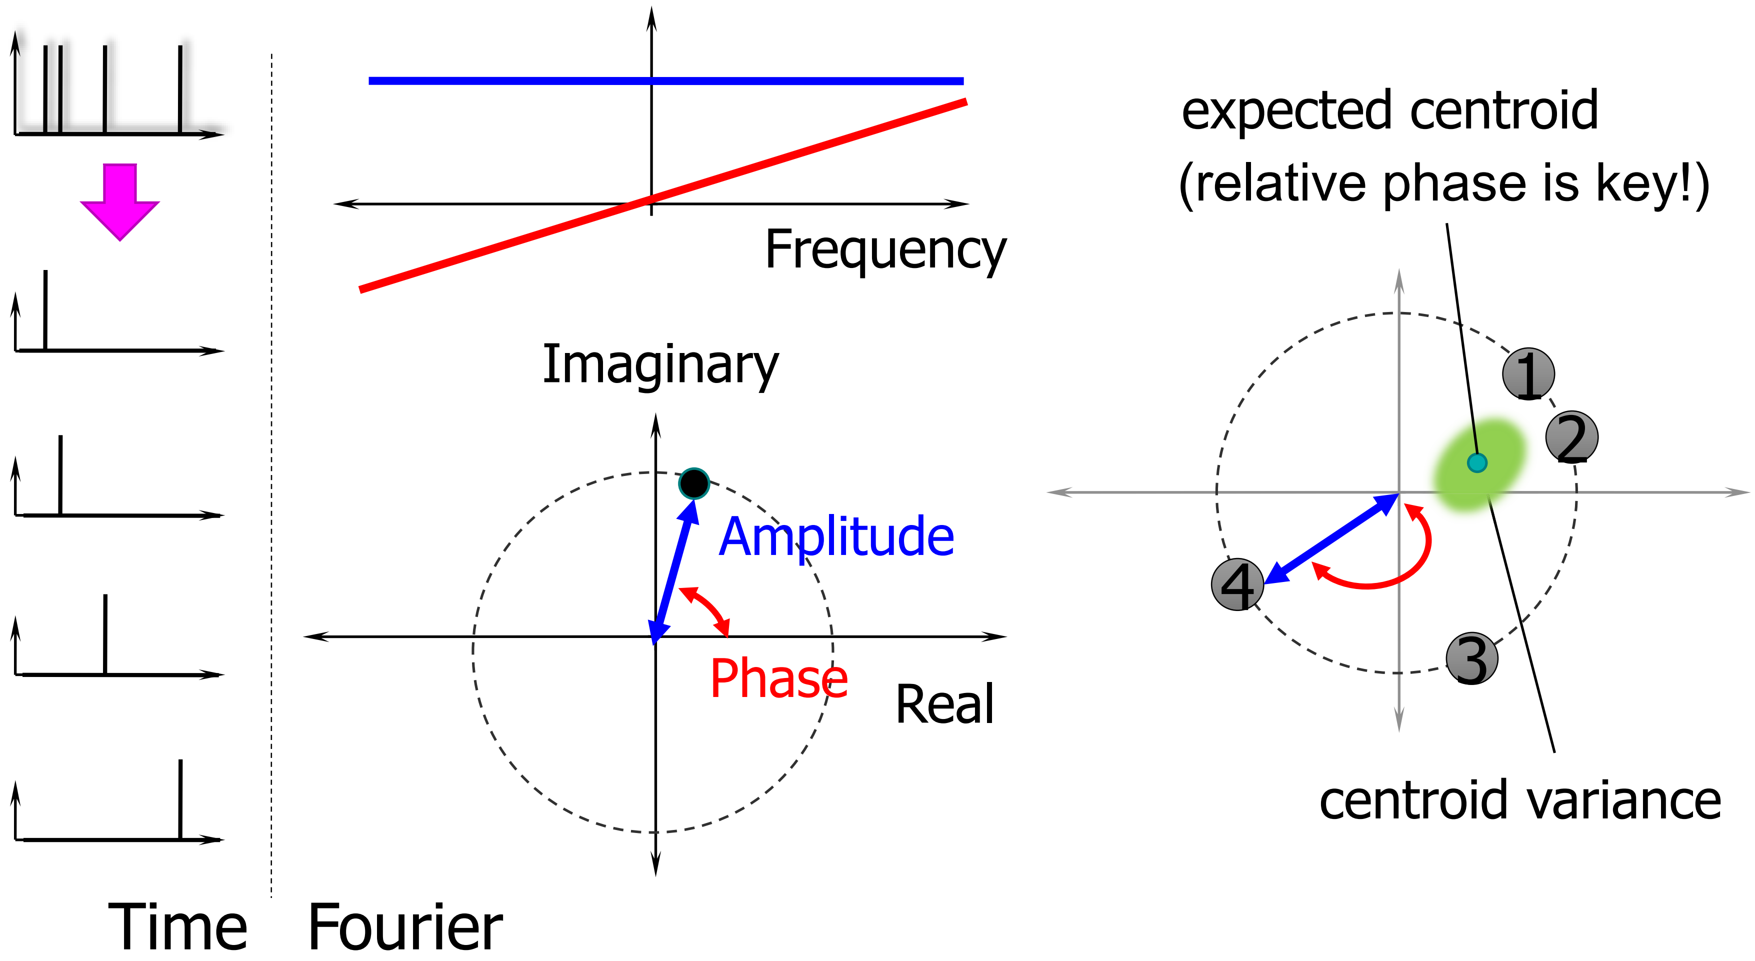

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

KAIST

Quantitative tests

- **Three types of synthetic datasets**

1) Four random samples to form a white quad (binary planes)

2) Delaunay triangulation of random samples (mesh) with random weights

3) Similar to Case 2. with linearly interpolated weights

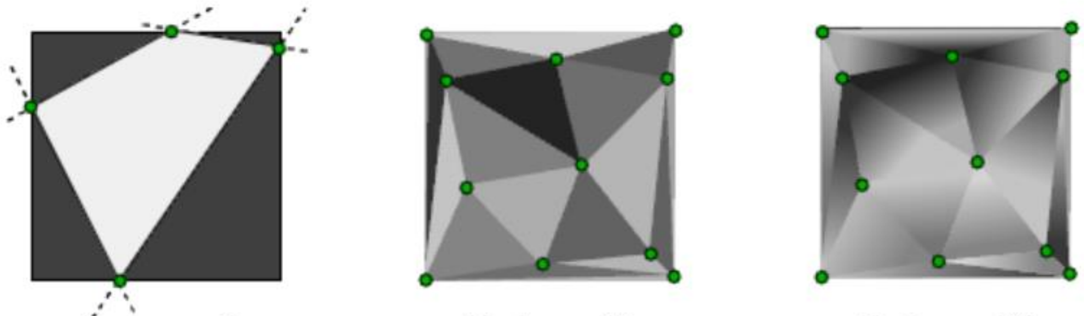

Test case I (binary) | Test case II (piece-wise constant) | Test case III (piece-wise linear)

- **Relative errors of mean, variance for methods in PBRT-v2 and other implem. (50 iterations of the secondary estimator with up to 1024 primary samples)**

Quantitative tests: Bias-variance trade-off using Gaussian jitter

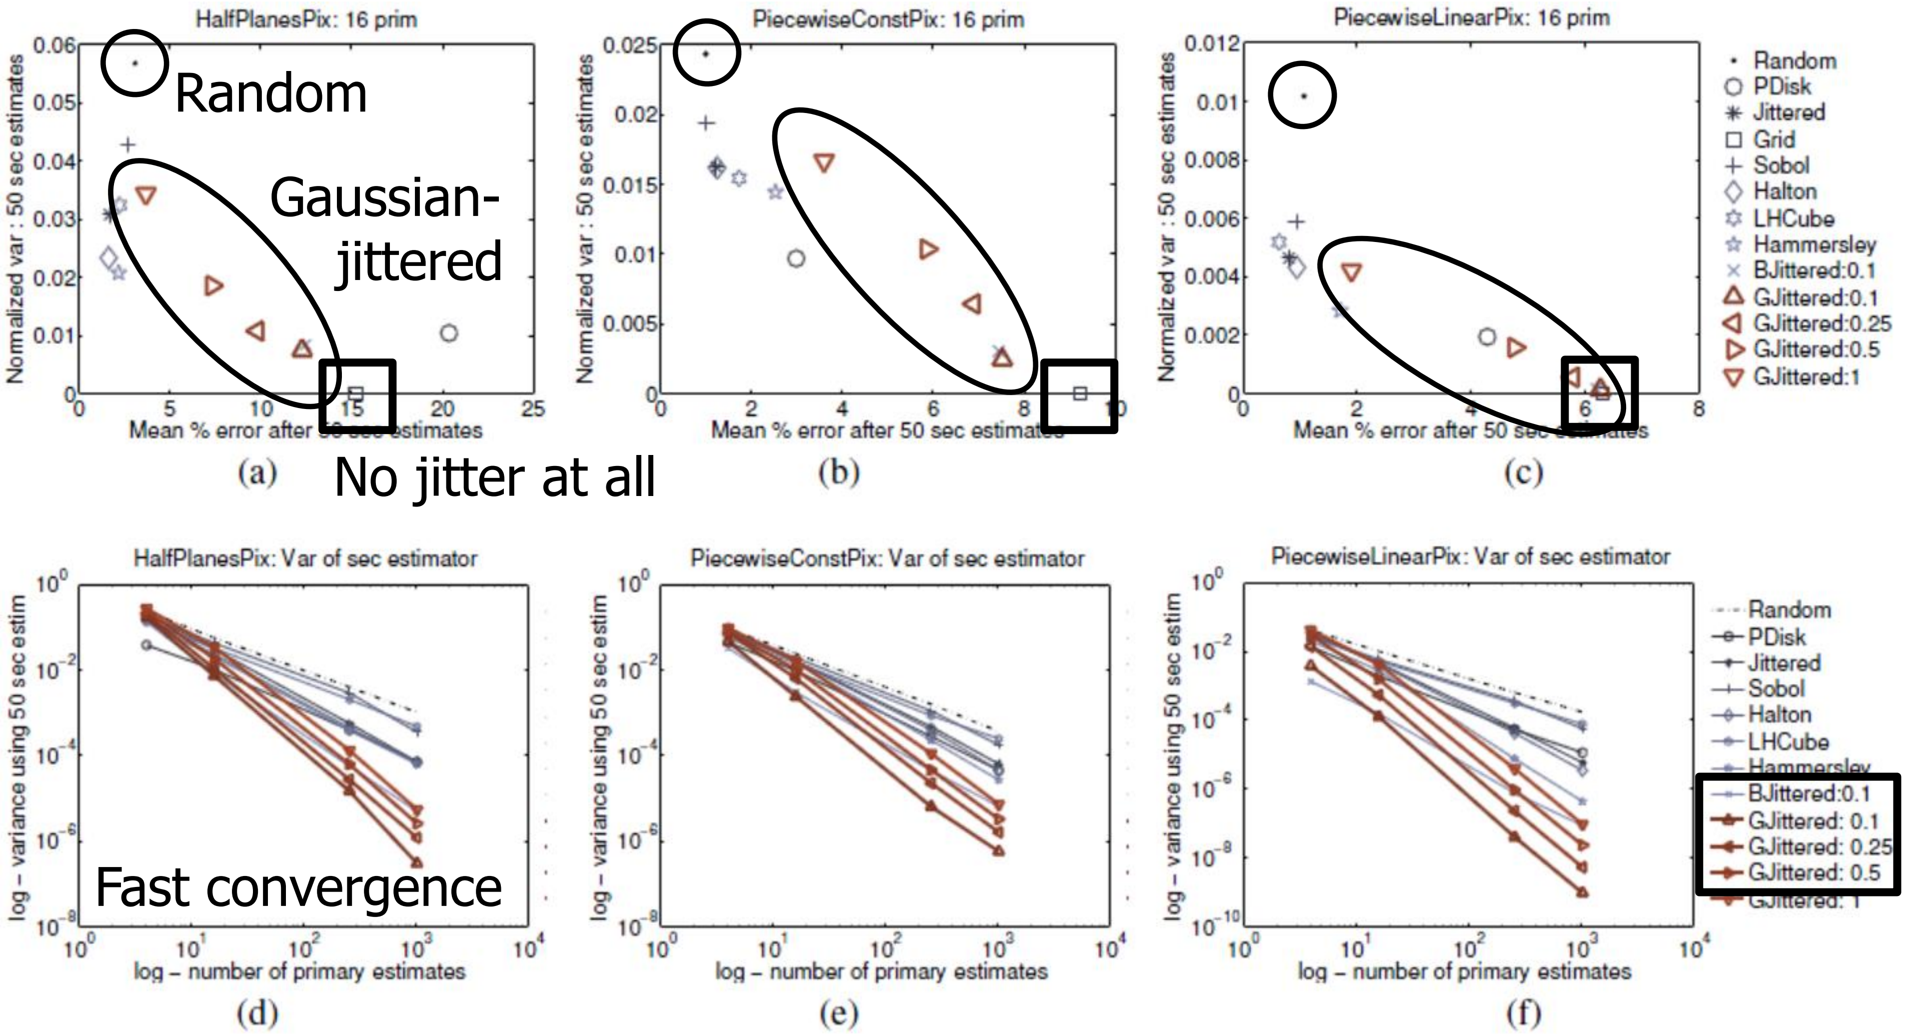

KAIST

Qualitative tests

- **Gaussian jitter allows trade-offs between ‘random’ and ‘grid.’**

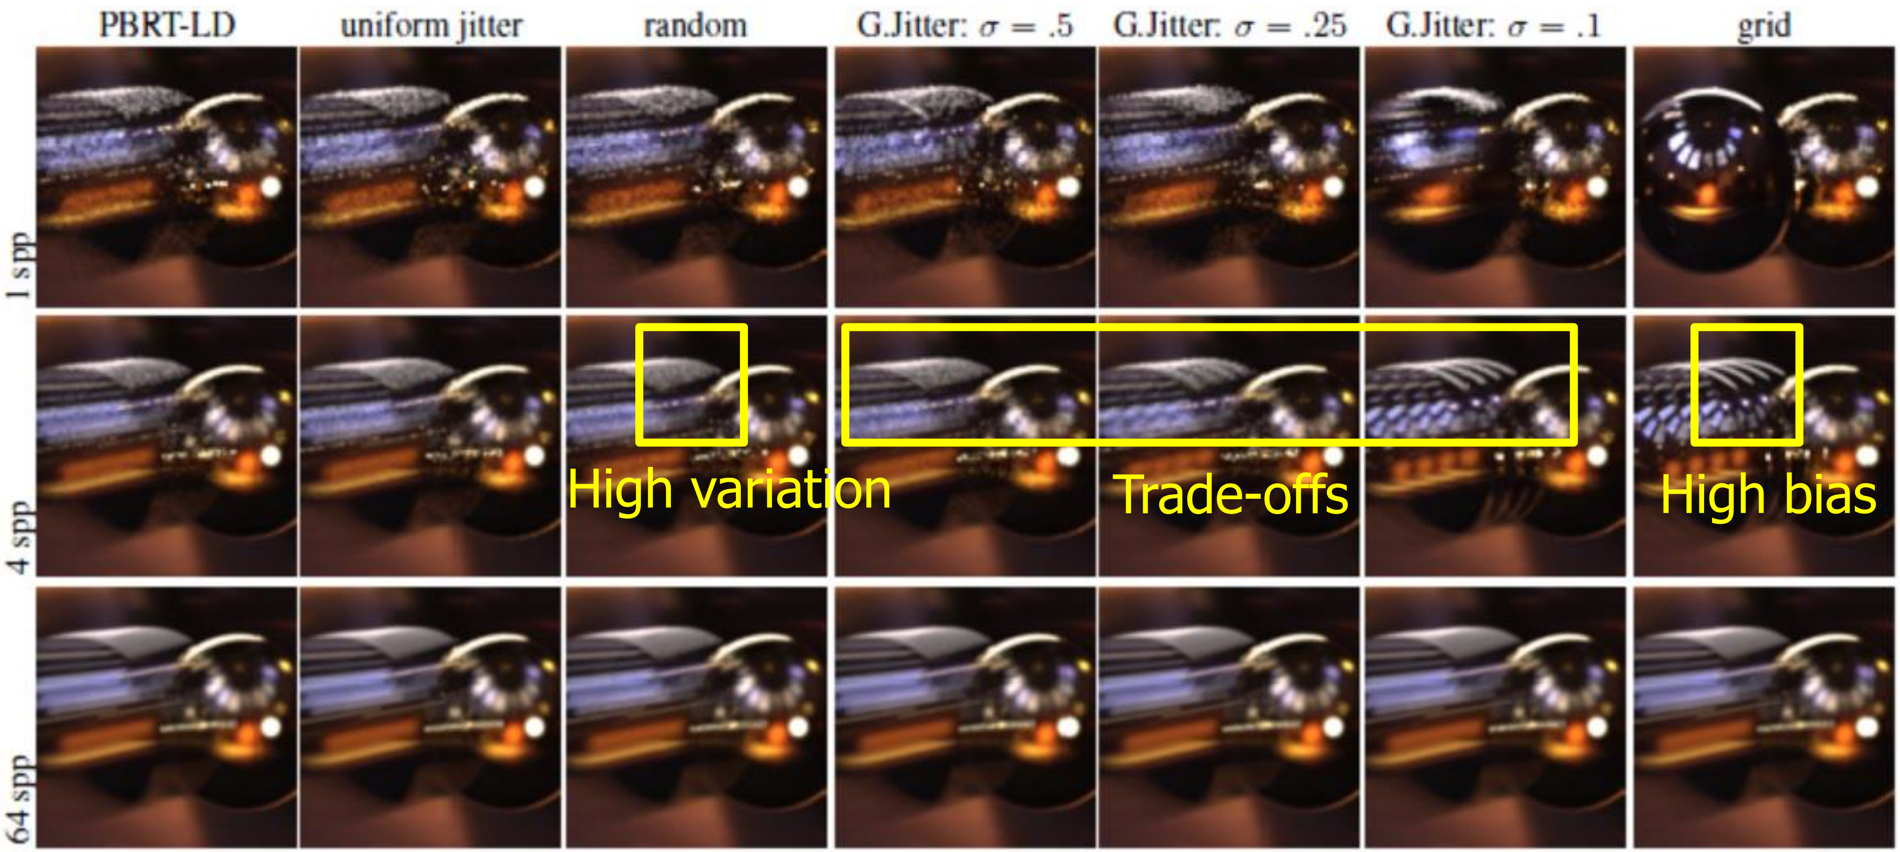

Conclusion

- **A study of the spectral characteristics of stochastic sampling patterns**
- **Two measures for the quality of sampling strategies in terms of their accuracy and precision in integration**
  - **The amplitude of the expected sampling spectrum**
  - **The variance of the sampling spectrum**
- **Applied these measures to assess Gaussian jittered sampling and compared it with the box-jittered case.**
- **Performed quantitative and qualitative evaluations of various sampling methods in this new framework.**

Appendix 1. Additional derivations

- **Gaussian jitter of stochastic samples: Expected spectrum**

$$\langle \hat{\mathbf{S}}'(\omega) \rangle \approx_2 \frac{1 - 2(\pi\omega\sigma)^2}{N} \langle \hat{\mathbf{S}}(\omega) \rangle$$

- **Gaussian jitter of fixed-location samples: Variance of spectrum**

$$\mathrm{V}\left(\mathbf{b}_\imath\right) \approx_2 (2\pi\omega\sigma)^2 \left(\cos^2(2\pi\omega\mathbf{X}_\imath) + \frac{(2\pi\omega\sigma)^2 \sin^2(2\pi\omega\mathbf{X}_\imath)}{2} \right)$$

similarly, the variance of the real part of $\langle \hat{\mathbf{S}}'(\omega) \rangle$

$$\mathrm{V}\left(\mathbf{a}_\imath\right) \approx_2 (2\pi\omega\sigma)^2 \left(\sin^2(2\pi\omega\mathbf{X}_\imath) + \frac{(2\pi\omega\sigma)^2 \cos^2(2\pi\omega\mathbf{X}_\imath)}{2} \right)$$

- **Spectrum for uniform jitter (1D):**

$$\begin{aligned} \langle \lim_{\omega \to k/T} \hat{\mathbf{S}}(\omega) \rangle &= \langle e^{-\imath 2\pi k \varsigma} \rangle \lim_{\omega \to k/T} \delta(\omega - k/T) \\ &= \begin{cases} \delta(\omega) & \text{if } k = 0 \\ 0 & \textit{otherwise} \end{cases} \end{aligned}$$

Bias and variance of secondary MC estimators

- **Statistics of Fourier spectrum over several sampling functions**

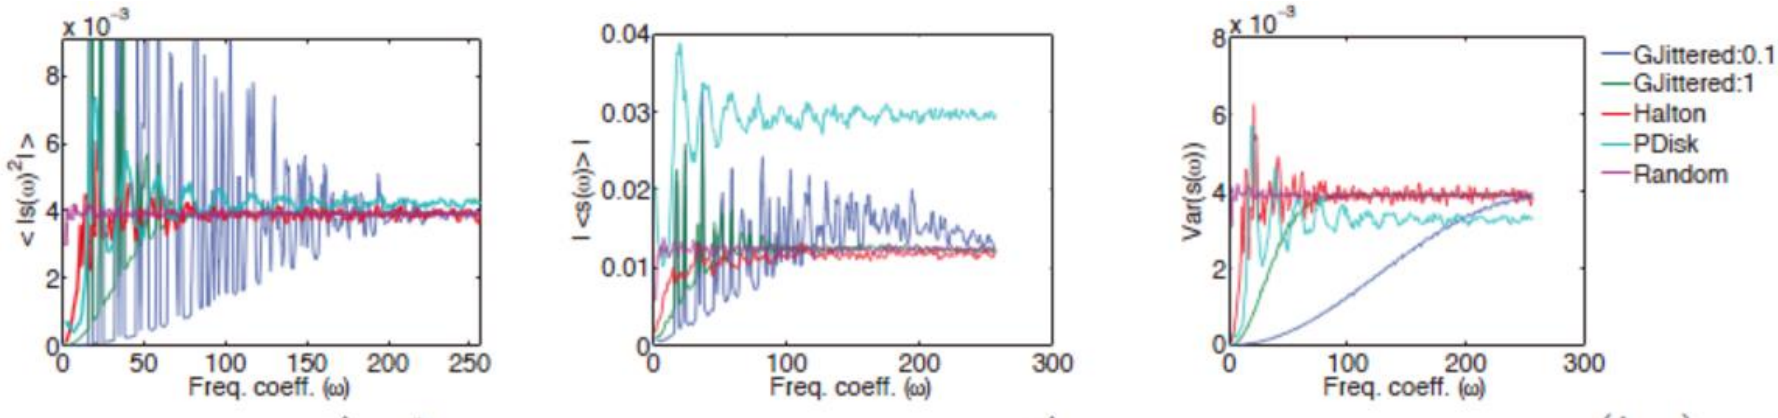

(a) Periodogram: $\langle |\hat{\mathbf{S}}(\omega)|^2 \rangle$ (b) Amplitude of expected spectrum: $|\langle \hat{\mathbf{S}}(\omega) \rangle|$ (c) Variance of spectrum: $\mathrm{V}\left(\hat{\mathbf{S}}(\omega)\right)$

- 512x512 grid, 256 2D samples, 20 iterations

- 5 sampling strategies (Gaussian jittered sampling, Halton sequences, Poisson-disk sampling, Random sampling)

- Amplitude of expected sampling spectrum, variance of sampling spectrum are more informative indicators for predicting stochastic sampling errors.

KAIST

Bias and variance of secondary MC estimators

- **Statistics of Fourier spectrum over several sampling functions**

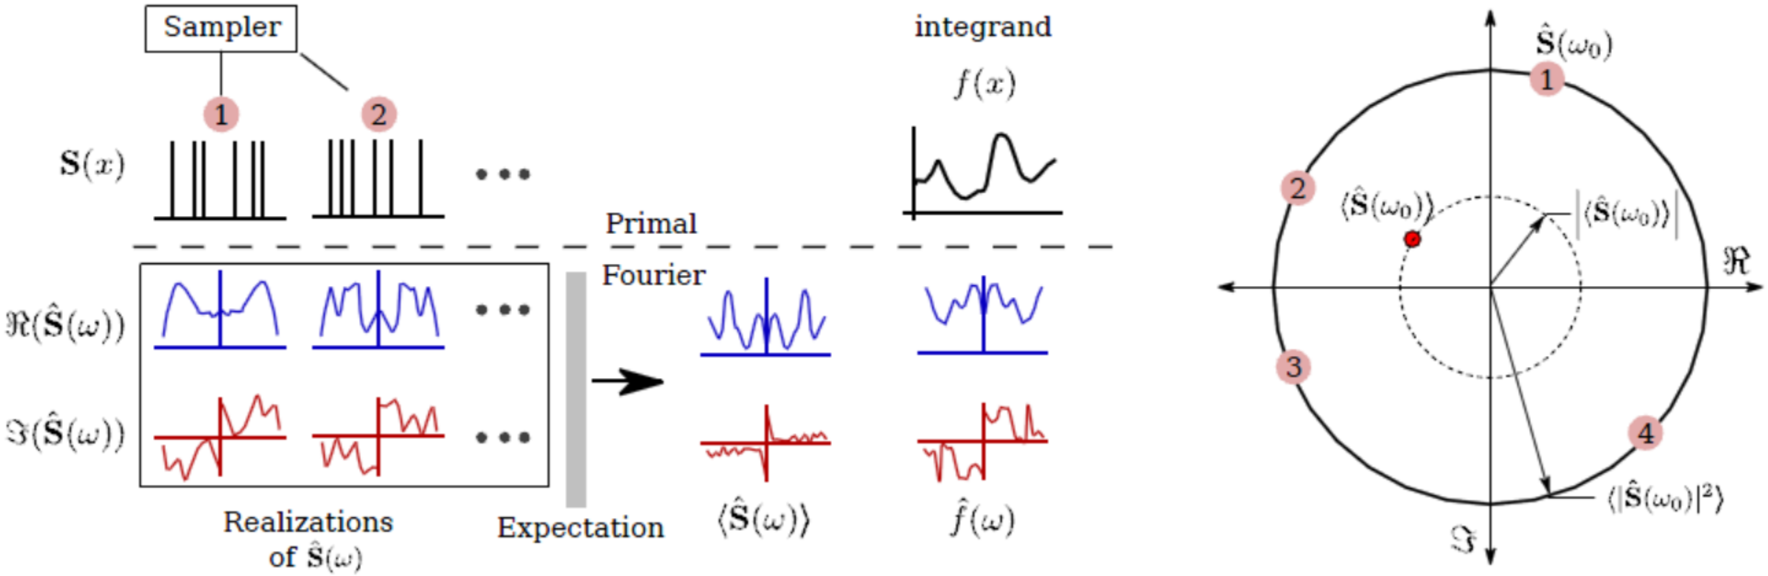

(a) Summary of notation

(b) Graphical interpretation of $|\langle \hat{\mathbf{S}}(\omega) \rangle|$

- Simple representations in each instance of the sampling patterns
(more expressive & informative than the conventional Periodogram)

* Shifted Dirac deltas are expressed as Euler formulas.

Case study: Gaussian jittered sampling

- **Expectation and variance of FT for Gaussian jittered samples**

$$\langle \hat{\mathrm{S}}'(\omega) \rangle \approx_2 \left(1 - \frac{(2\pi\omega\sigma)^2}{2}\right) \frac{\hat{\text{Ш}}_N(\omega)}{N}$$

$$\mathrm{V}\left(\hat{\mathrm{S}}'(\omega)\right) \approx_2 \frac{(2\pi\omega\sigma)^2}{N}(1 + 2(\pi\omega\sigma)^2)$$

- **Effect of the jitter parameter σ in terms of the frequency**

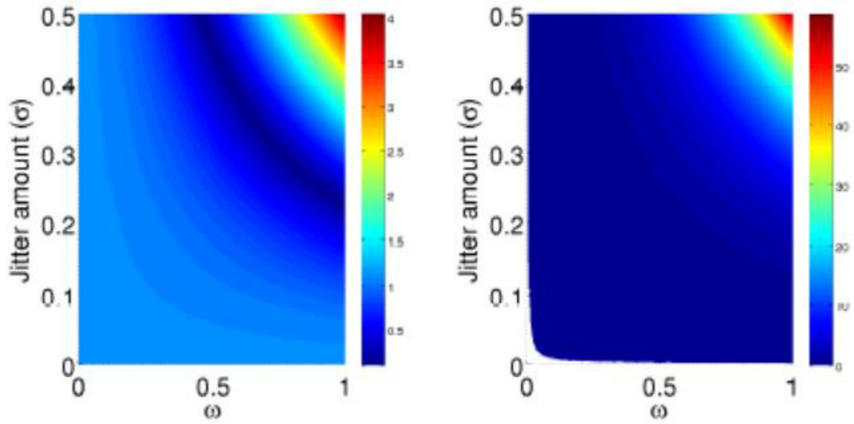

(a) $|\langle \hat{\mathrm{S}}'(\omega) \rangle : \langle \hat{\mathrm{S}}(\omega) \rangle|$ (b) $\mathrm{V}\left(\hat{\mathrm{S}}(\omega)\right)$

- No bias for pure DC signals.
- Jitter increases bias beyond yellow regions. (high freq.)
- σ should be small enough to fall in the blue regions.

KAIST

Case study: Gaussian jittered sampling

- **Expectation and variance of FT for Box-jittered samples**

$$\langle \hat{\mathrm{S}}'(\omega) \rangle \approx_2 \left(1 - \frac{(2\pi\omega\sigma)^2}{6}\right) \frac{\hat{Ш}_N(\omega)}{N}$$

$$\mathrm{V}\left(\hat{\mathrm{S}}'(\omega)\right) \approx_2 \frac{(2\pi\omega\sigma)^2}{N} \left(\frac{1}{3} + \frac{4}{45}(\pi\omega\sigma)^2\right)$$

- **Comparison of box-jitter and Gaussian jitter**

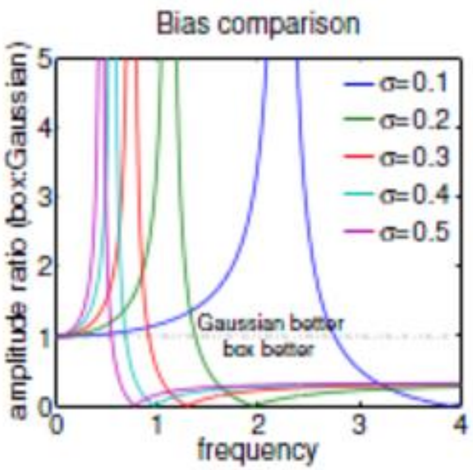

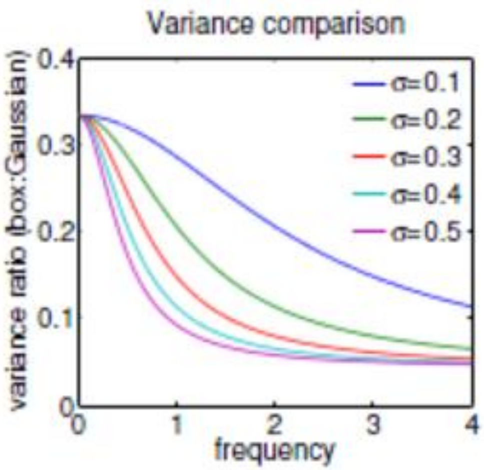

- Simple alternative to Gaussian
- Box jitter is more biased if the integrand contains energy at frequencies where the ratio is greater than 1.
- Variance of box filter is lower.

Quantitative tests: Gaussian jitter converges rapidly

	I: binary	II: p/w const	III: p/w linear
Random	$O(N^{-1.0020})$	$O(N^{-0.9745})$	$O(N^{-1.0098})$
Poisson disk	$O(N^{-1.0605})$	$O(N^{-1.1902})$	$O(N^{-1.3629})$
Jittered	$O(N^{-1.4171})$	$O(N^{-1.1877})$	$O(N^{-1.5003})$
Sobol	$O(N^{-1.0200})$	$O(N^{-1.0374})$	$O(N^{-1.0879})$
Halton	$O(N^{-1.4233})$	$O(N^{-1.2927})$	$O(N^{-1.5890})$
LHCube	$O(N^{-1.0112})$	$O(N^{-1.1255})$	$O(N^{-1.0903})$
Hammersley	$O(N^{-1.3813})$	$O(N^{-1.2324})$	$O(N^{-1.7772})$
BJittered:0.1	$O(N^{-1.8969})$	$O(N^{-1.6161})$	$O(N^{-1.7370})$
GJittered:0.1	$O(N^{-2.2478})$	$O(N^{-2.1373})$	$O(N^{-2.6781})$
GJittered:0.25	$O(N^{-2.1368})$	$O(N^{-1.9015})$	$O(N^{-2.5759})$
GJittered:0.5	$O(N^{-1,9464})$	$O(N^{-1.7070})$	$O(N^{-2.3969})$
GJittered:1	$O(N^{-1.7542})$	$O(N^{-1.5311})$	$O(N^{-2.0955})$

Qualitative tests: Blur, Soft shadow

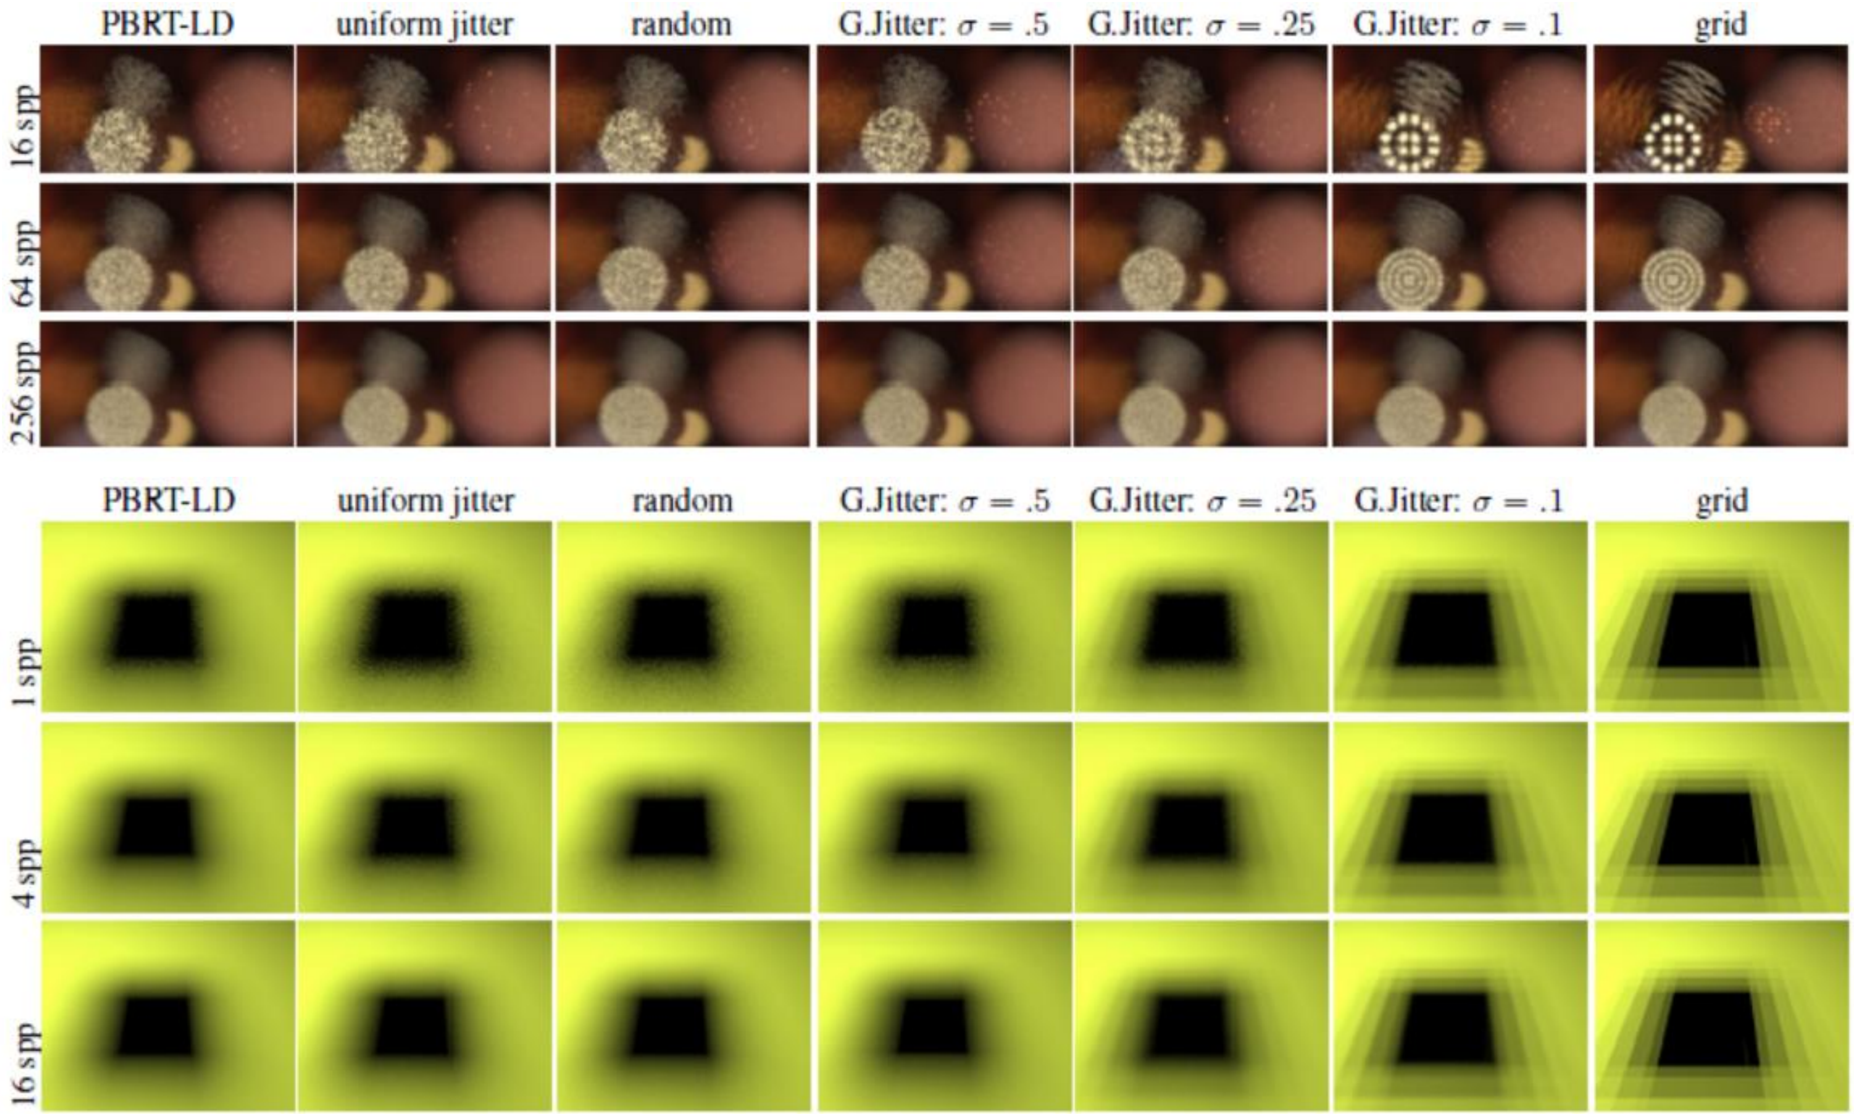

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views

Wants to know this viewpoint from a photo → 3D Viewpoint Estimation

(One of classical CV problems)

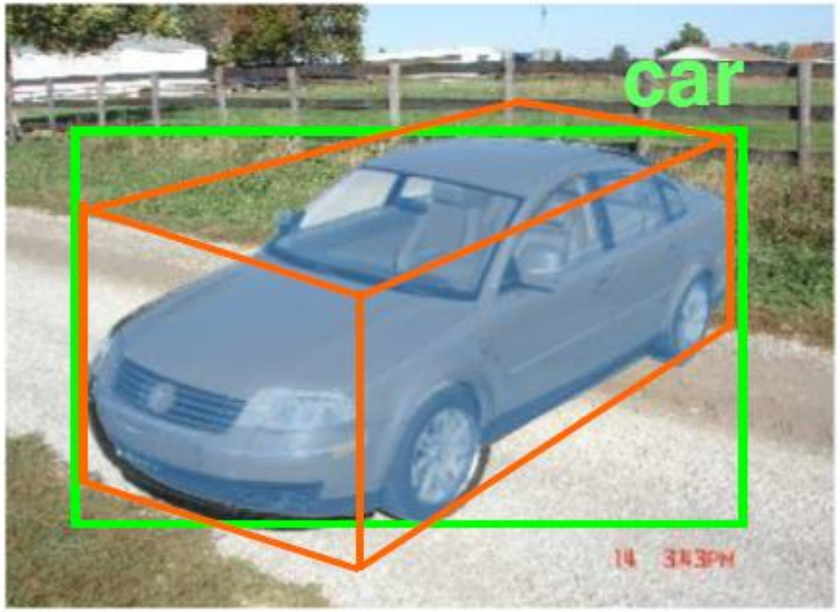

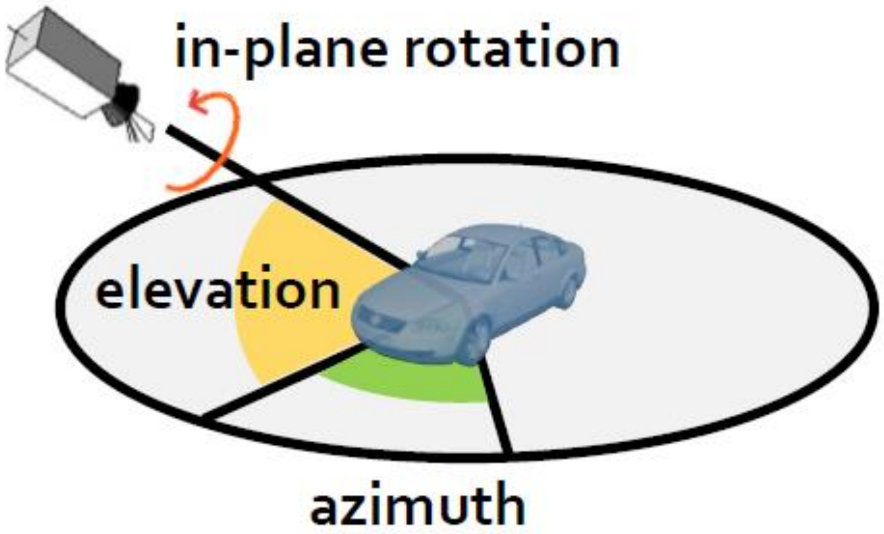

- **If 2D renders = 2D photos, our high-level descriptions naturally leads to high-level understandings for machines.**
- **For example, let's take a look at estimating camera viewpoints.**

Fourier Analysis of Stochastic Sampling Strategies for Assessing Bias and Variance in Integration
Author's Slides from http://home.eps.hw.ac.uk/~ks400/research.html

Convolutional Neural Network (CNN)

- **ImageNet: Millions of images + Human annotations (2009)**
- **ILSVRC Image Classification Top-5 Error (%)**

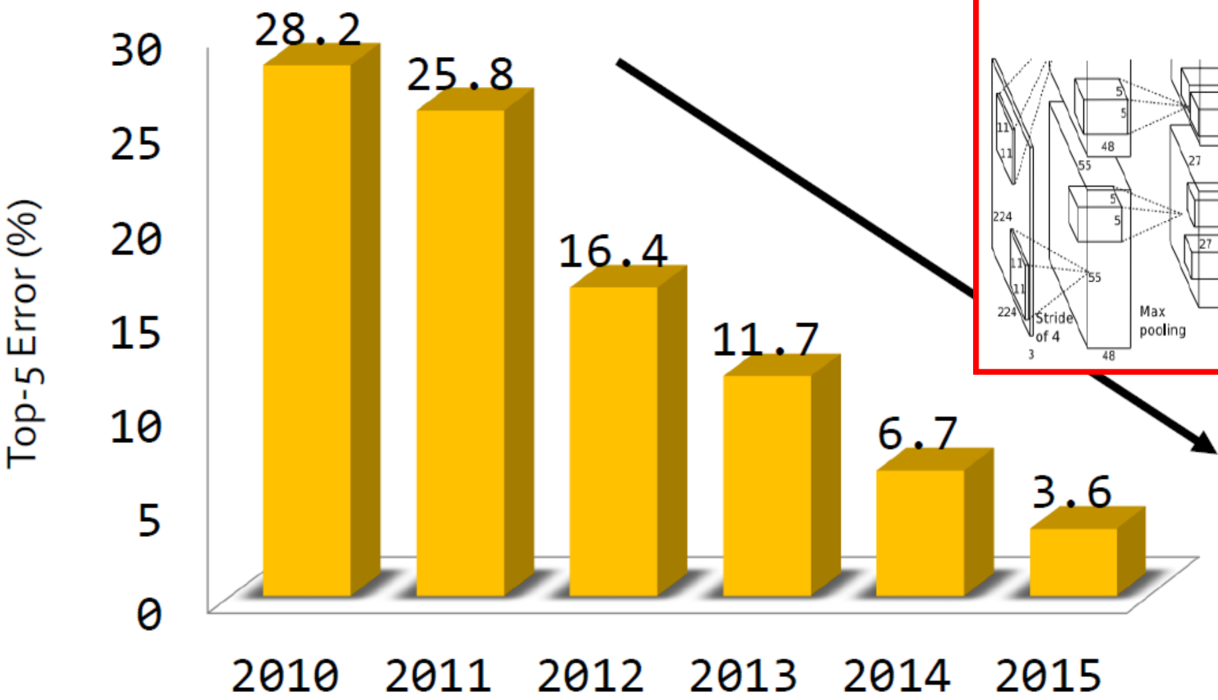

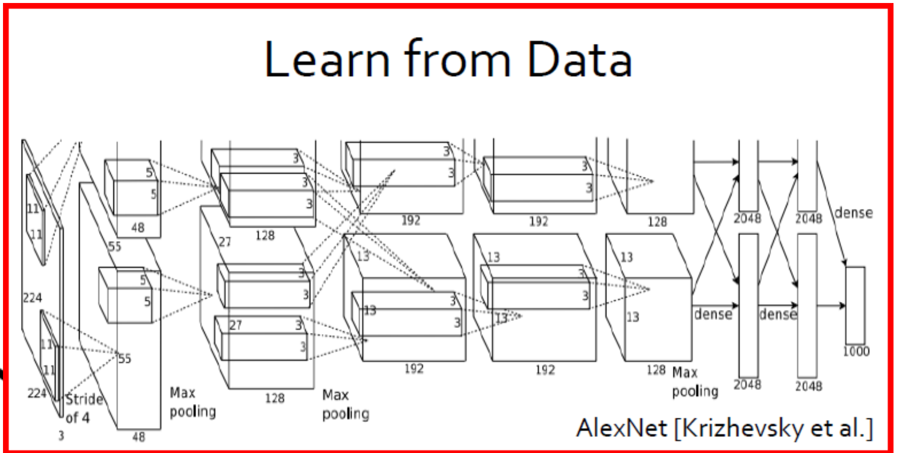

CPU↑ GPU↑

Deep Learning!

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

KAIST

Introduction

- **Go beyond 2D image classification**

What's the camera viewpoint angles to the SUV in the image?

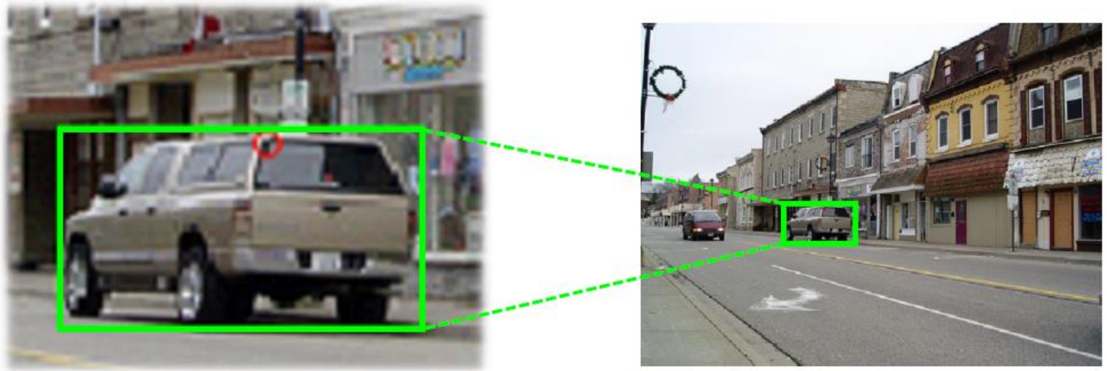

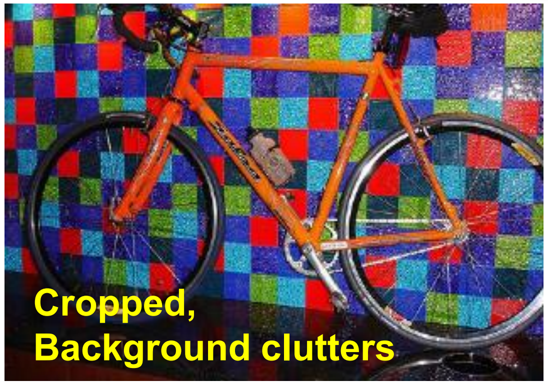

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Introduction

- **Our challenge here is that human annotation is expensive.**

PASCAL3D+ dataset [Xiang et al.]

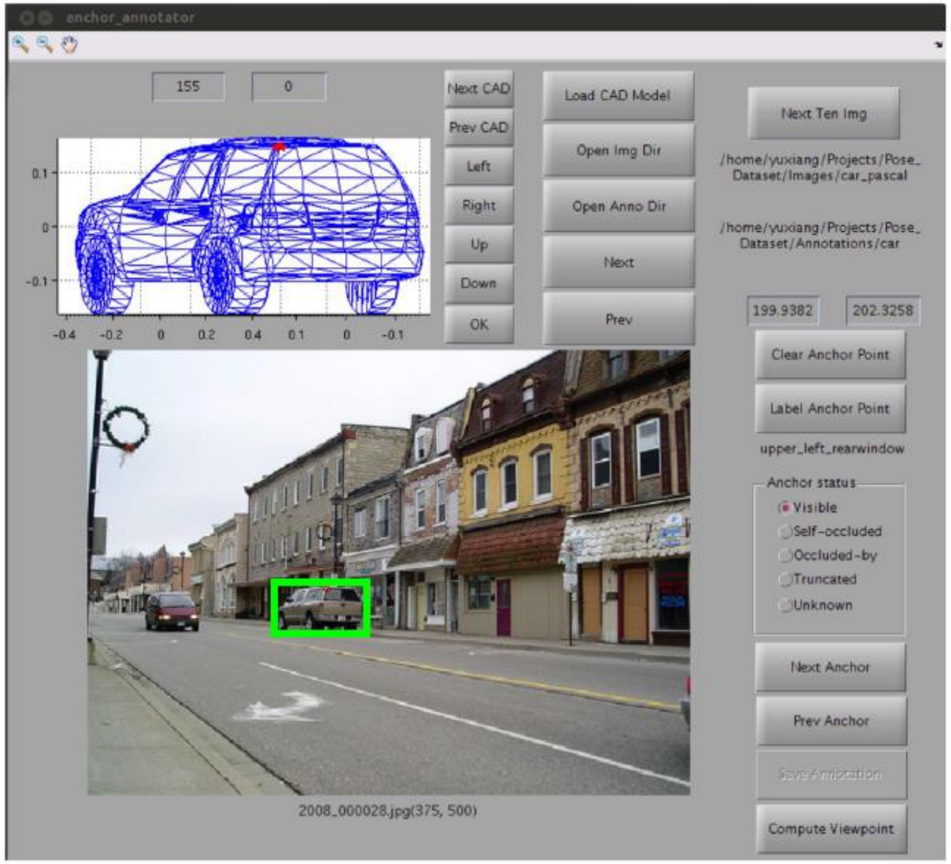

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Introduction

High-cost Label Acquisition

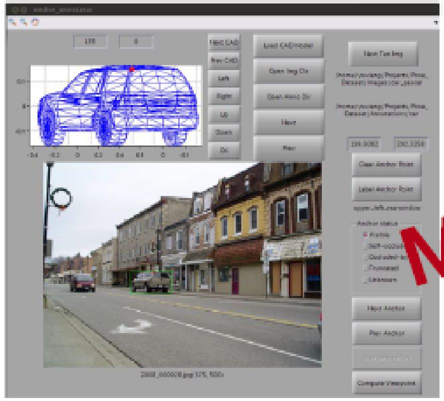

30K images with viewpoint labels in PASCAL3D+ dataset [Xiang et al.]

Mismatch!

High-capacity Model

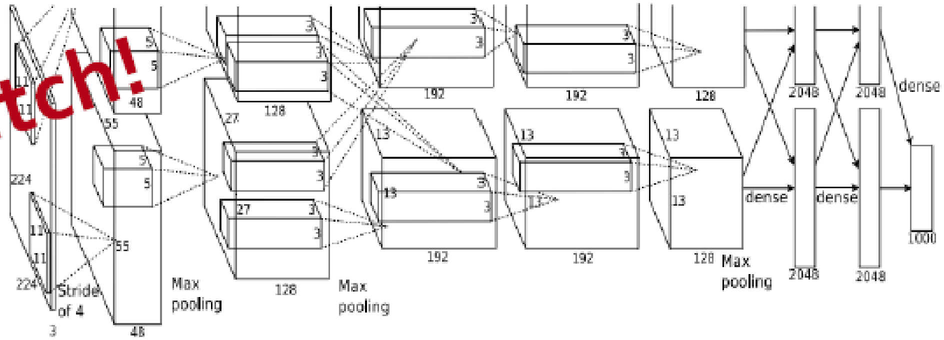

60M parameters. AlexNet [Krizhevsky et al.]

How to get MORE images with ACCURATE viewpoint labels?

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)
KAIST

Key idea: Render for CNN

- **CG renderings are generated by known model descriptions**
  ***All data is already annotated when created.**

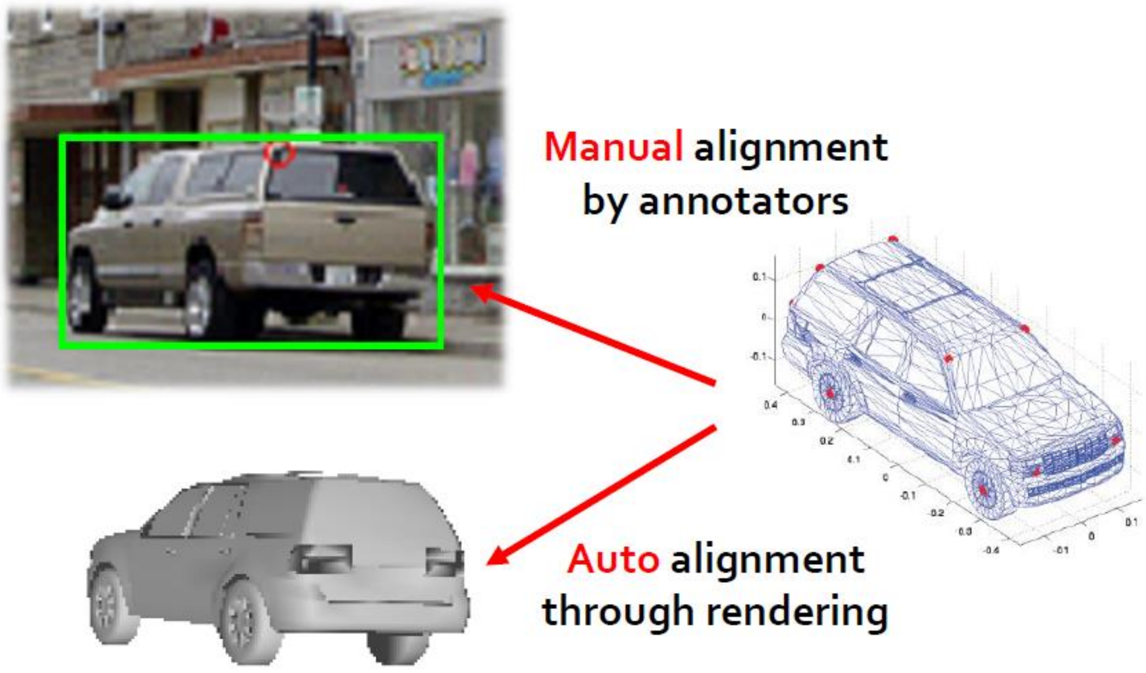

Renderings
from 3D models

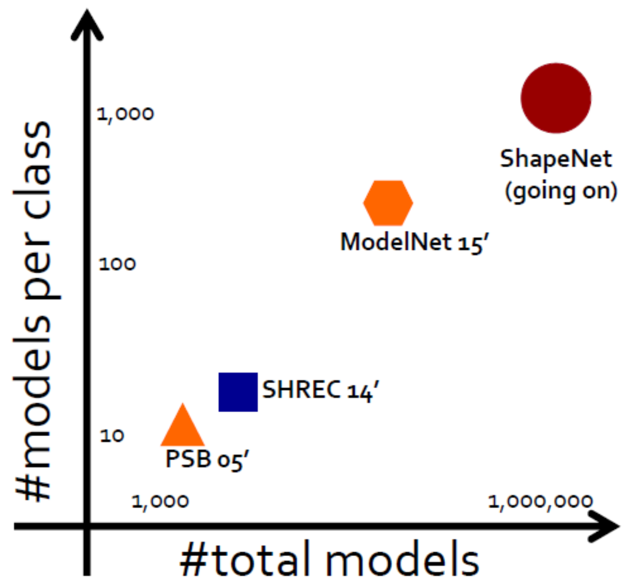

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Rendering pipeline for the training stage**

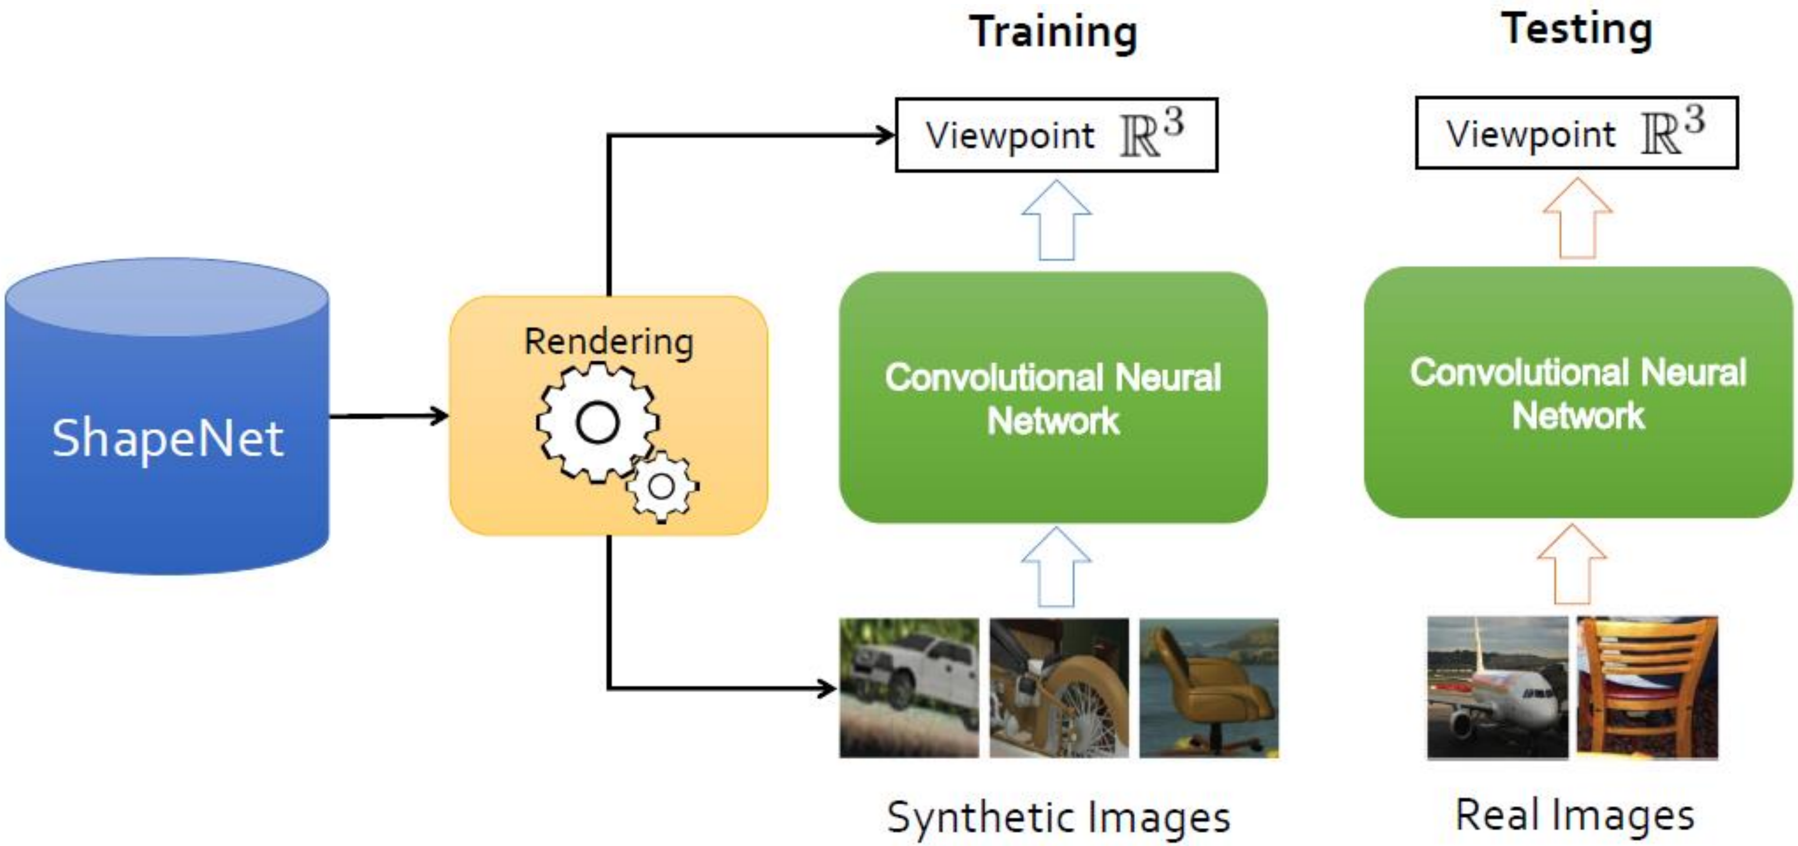

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Self-generated data collections for machine learning**

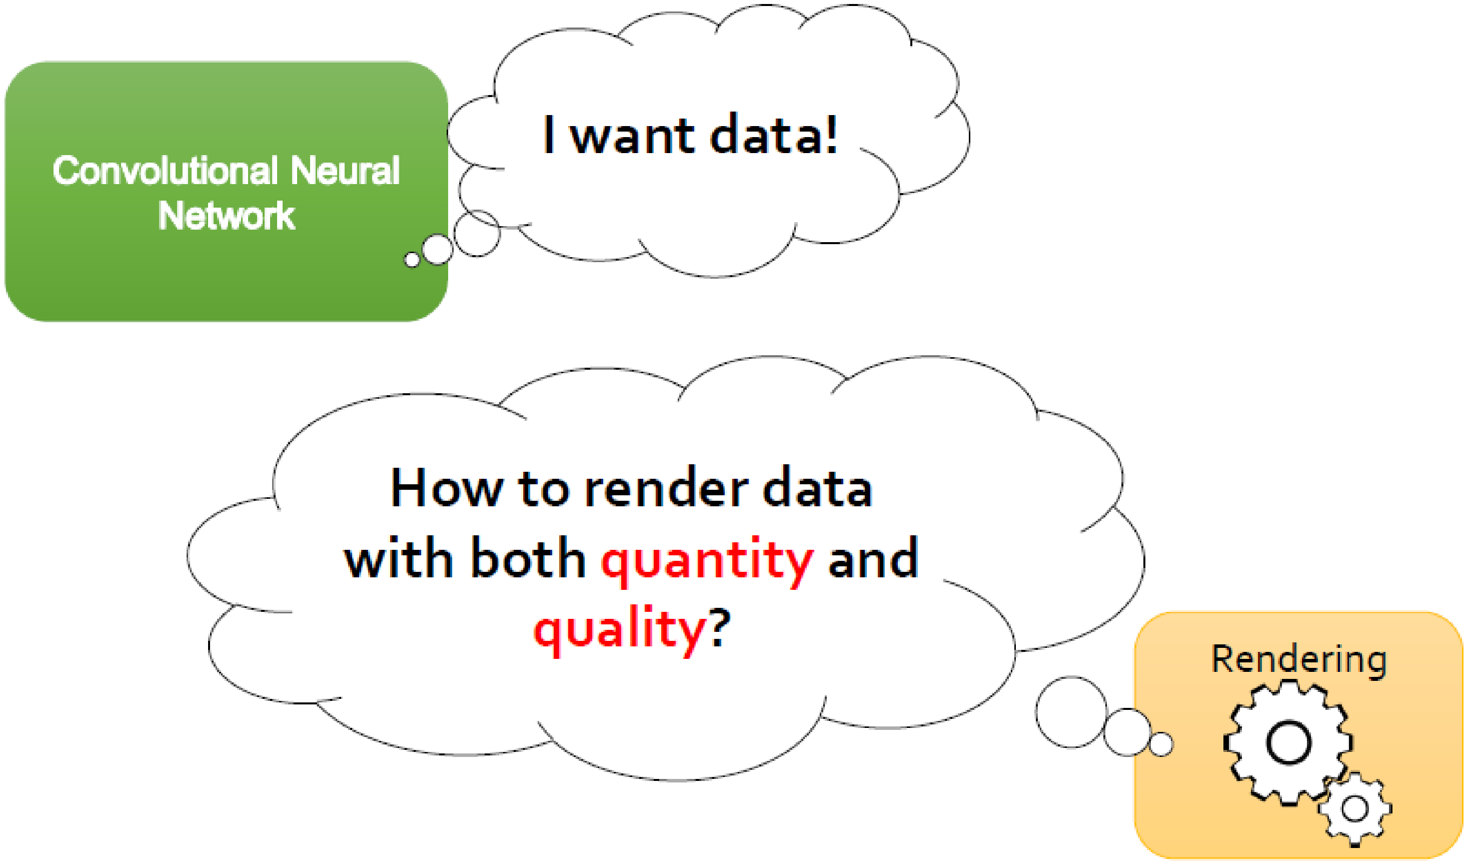

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Exp. 1. 80K rendered chair images with fixed lighting sources**

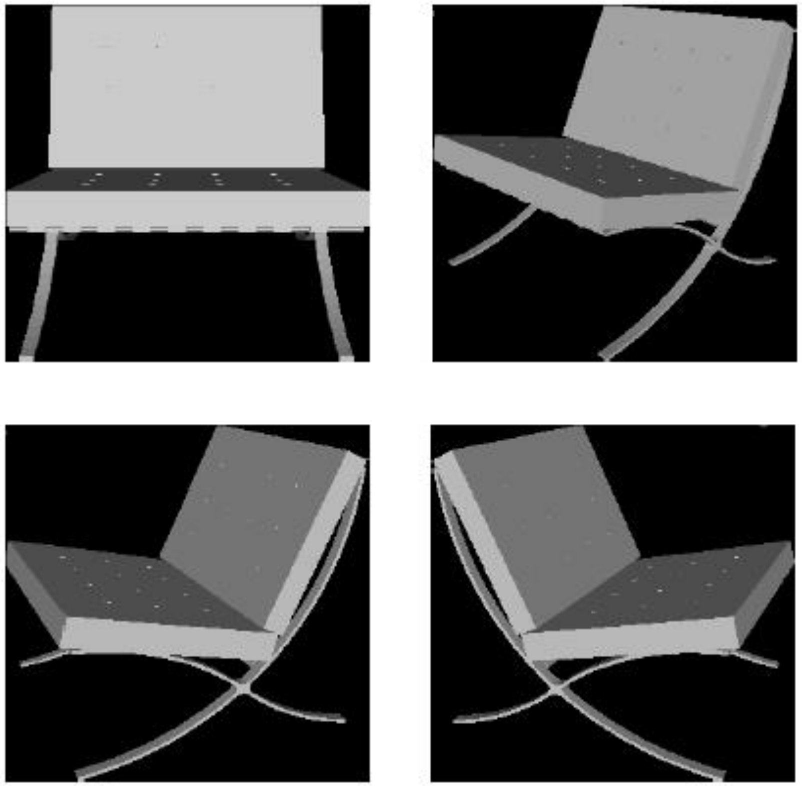

95% on synthetic val set

47% on real test set ☹

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Exp. 2. Randomize lighting**

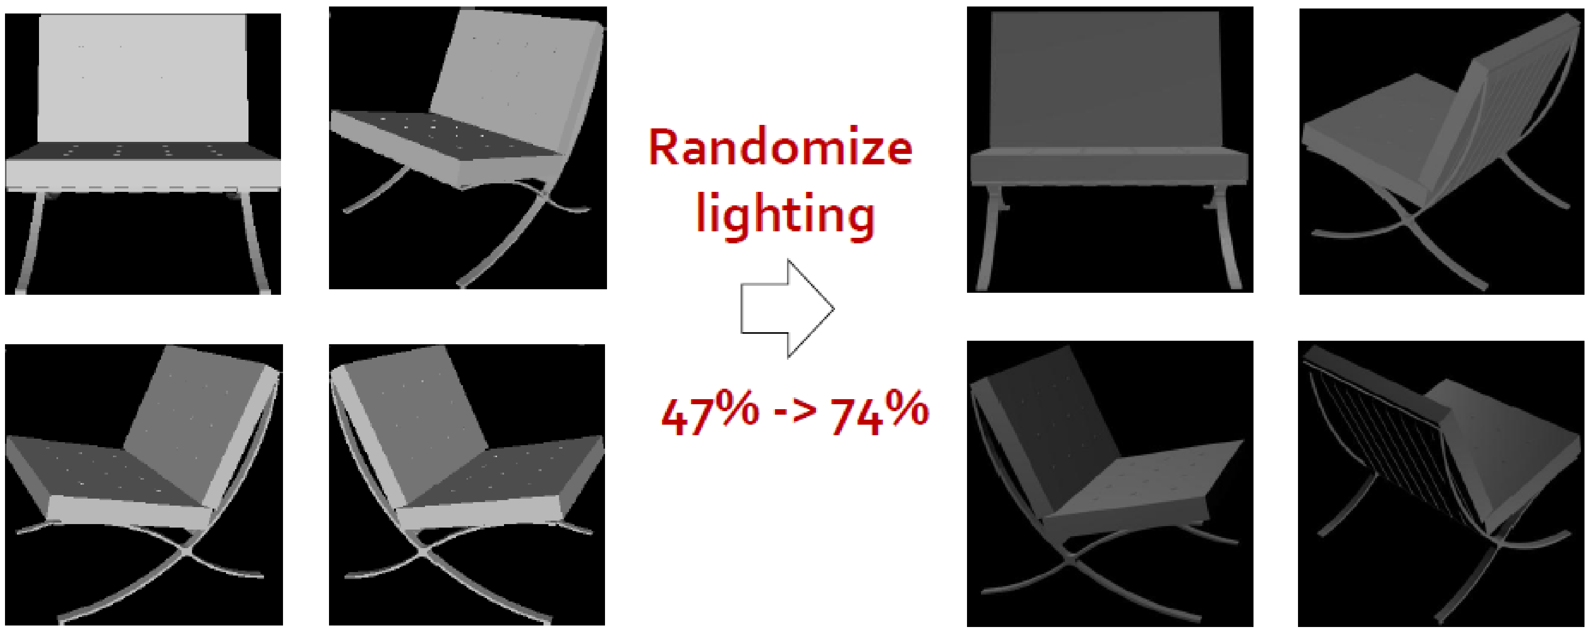

ConvNet: hmm.. viewpoint is not the brightness pattern. Maybe it's the contour?

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Exp. 3. Composite them with random backgrounds**

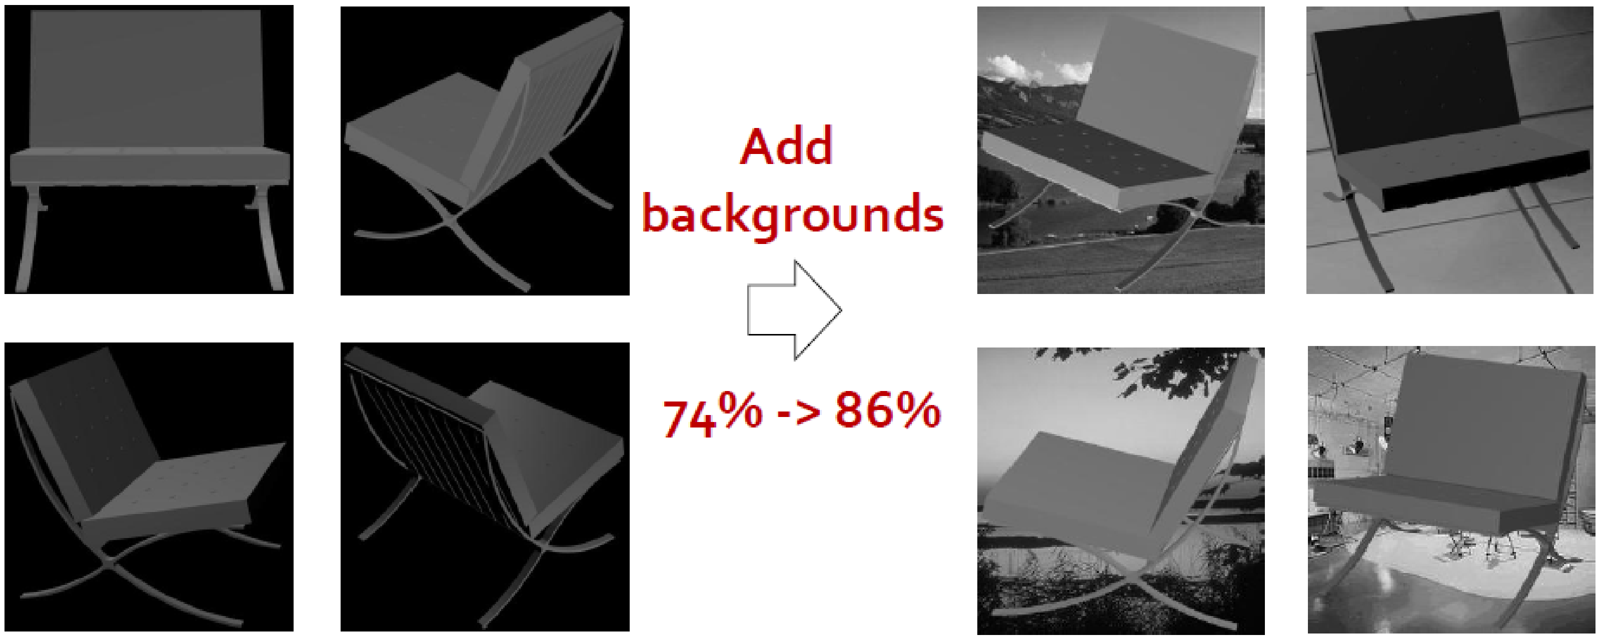

ConvNet: It becomes really hard! Let me look more into the picture.

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **Exp. 4. Apply bounding boxes with proper texture**

bbox crop
texture

86% -> 93%

Key Lesson: Don't give CNN a chance to "cheat" - it's very good at it. When there is no way to cheat, true learning starts.

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Rendering pipeline

- **4M synthesized images for 12 categories**

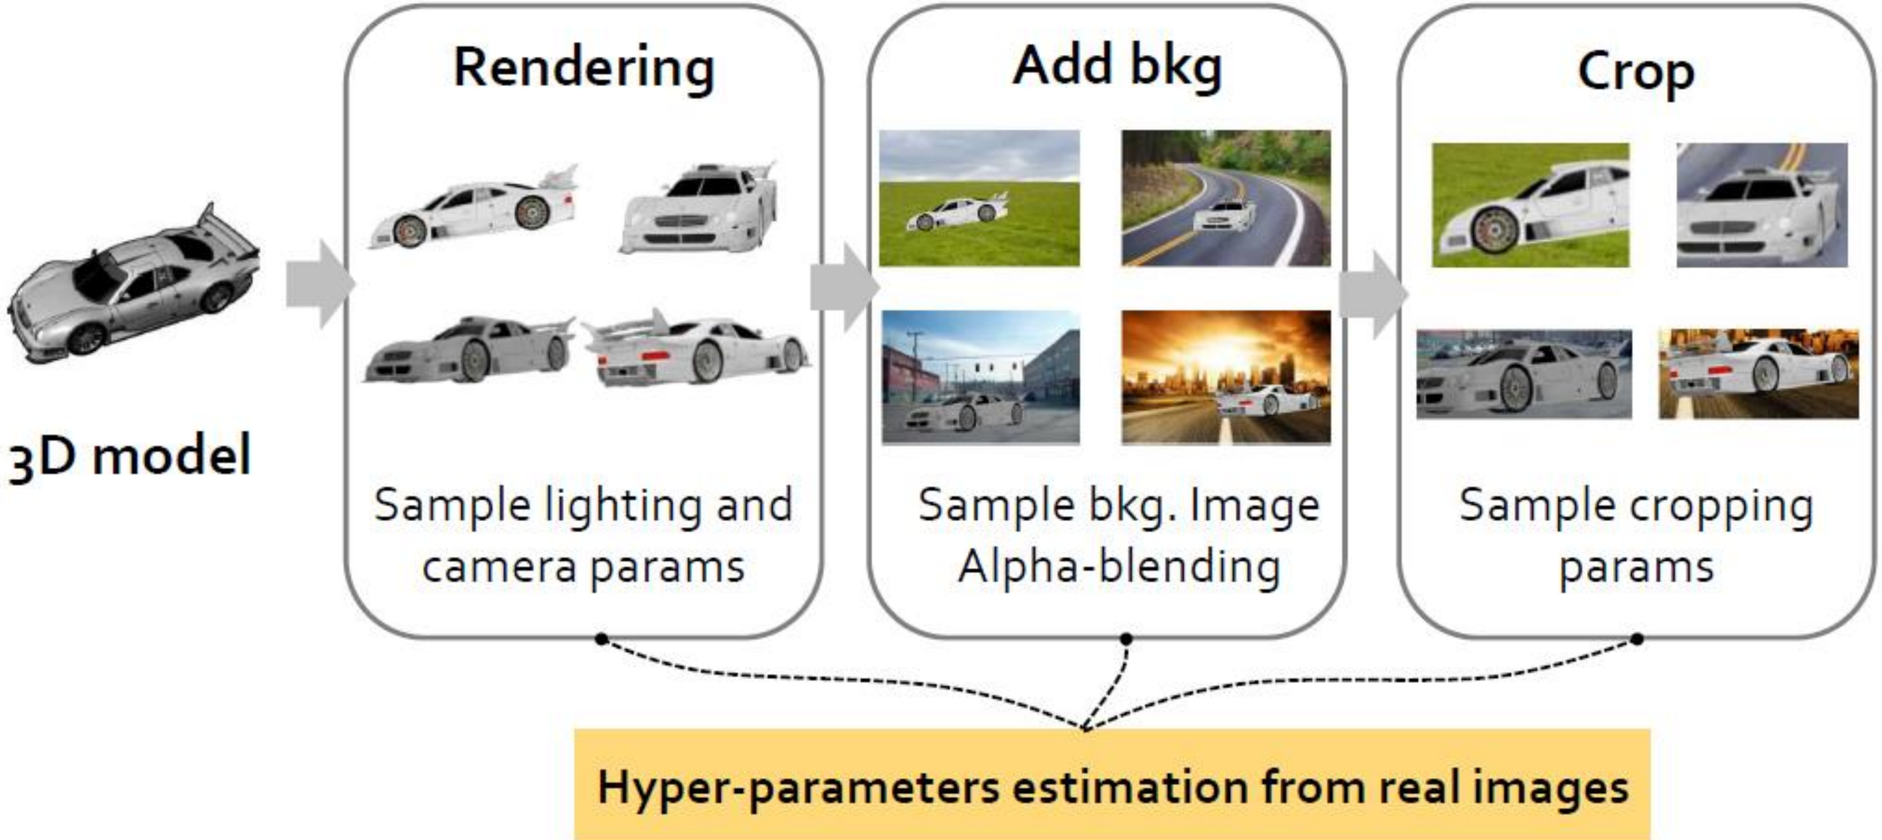

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Experimental results

- **Real test images from PASCAL3D+ dataset**
- **Metric: median angle error (lower the better)**

	aero	bike	boat	bottle	bus	car	chair	table	mbike	sofa	train	tv	mean
$Acc_{\frac{\pi}{6}}$ (Tulsiani, Malik)	0.78	0.74	0.49	0.93	0.94	0.90	0.65	0.67	0.83	0.67	0.79	0.76	0.76
$Acc_{\frac{\pi}{6}}$ (Ours-Render)	0.74	0.83	0.52	0.91	0.91	0.88	0.86	0.73	0.78	0.90	0.86	0.92	0.82
$MedErr$ (Tulsiani, Malik)	14.7	18.6	31.2	13.5	6.3	8.8	17.7	17.4	17.6	15.1	8.9	17.8	15.6
$MedErr$ (Ours-Render)	15.4	14.8	25.6	9.3	3.6	6.0	9.7	10.8	16.7	9.5	6.1	12.6	11.7

Our model **trained on rendered images** outperforms state-of-the-art model **trained on real images** in PASCAL3D+.

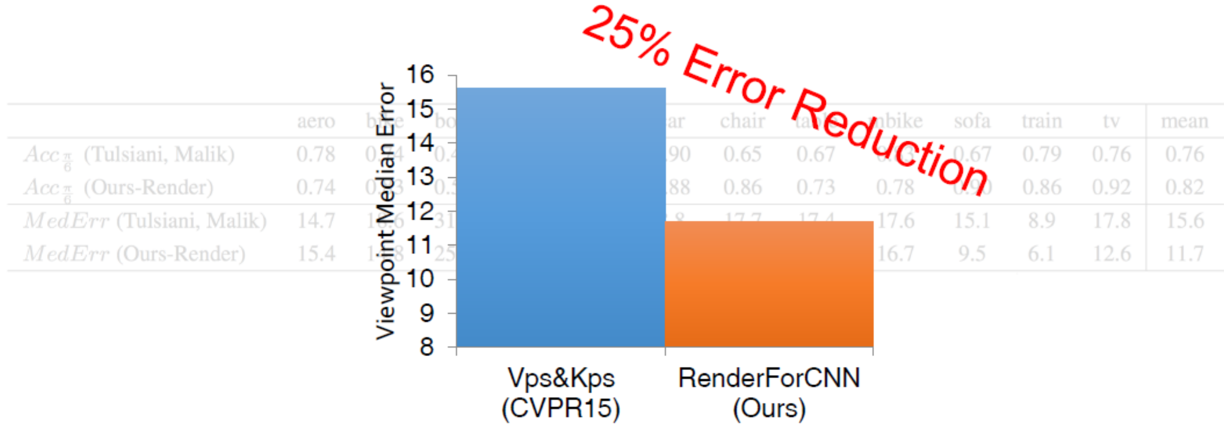

Experimental results

- **Azimuth Viewpoint Estimation**

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Experimental results

- **Failure cases**

sofa occluded by people

ambiguous car viewpoint

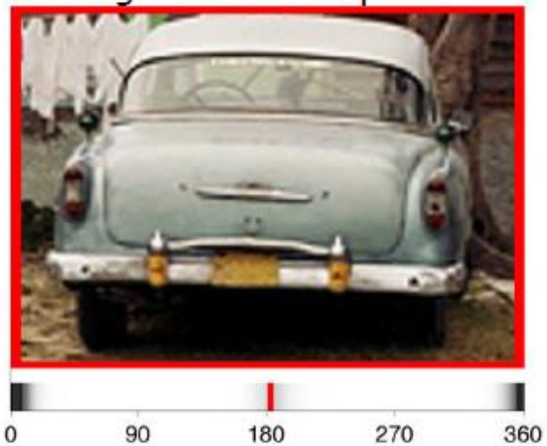

multiple cars

car occluded by motorbike

ambiguous chair viewpoint

multiple chairs

Conclusion

- **Images rendered from 3D models can be effectively used to train CNNs, especially for 3D tasks. State-of-the-art results has been achieved.**

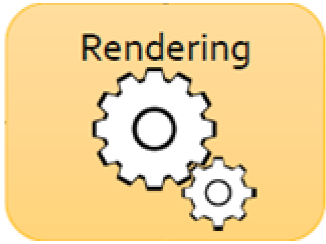

*** Self-imaging from descriptions**

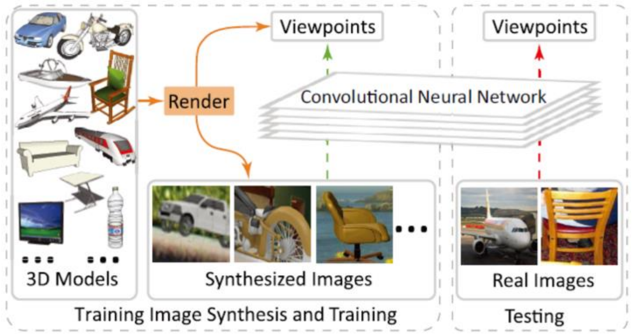

- **Key to success**
  - **Quantity: Large scale 3D model collection (ShapeNet) + Augmentation**
  - **Quality: Overfit-resistant, scalable image synthesis pipeline**

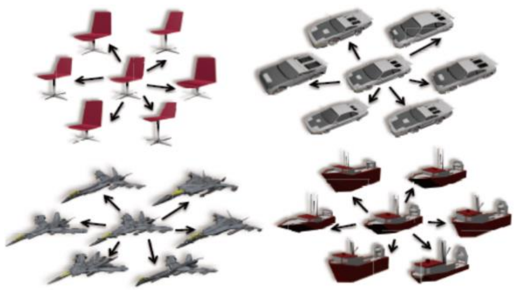

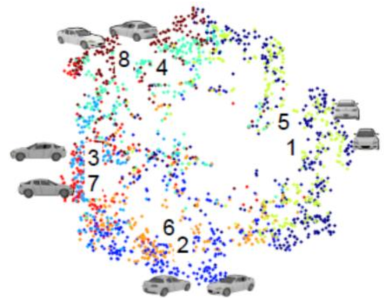

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Appendix 2. Convolutional Neural Network (CNN)

- **Multiple layers of small neuron collections**

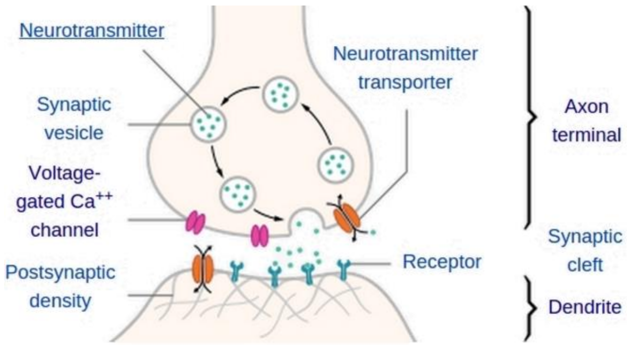

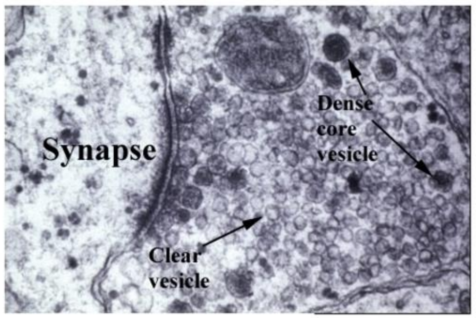

*** Parametric model for simulating a neural network:**

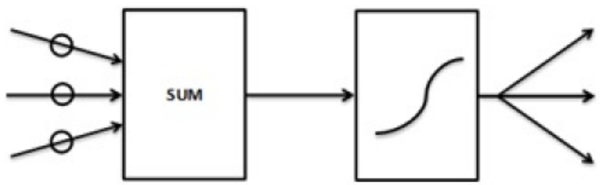

Weighted inputs (Parameters) | **Limiter (non-linear step-function)** | **Output to Other neurons**

*** Application: Image classification→ any semantic inferences**

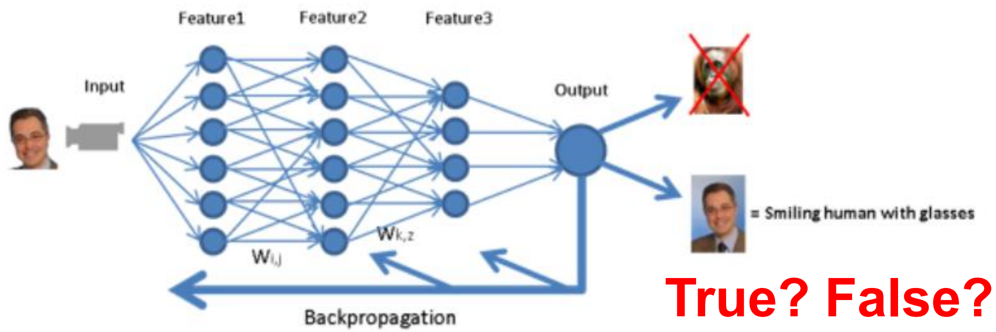

True? False?

Require a huge number of labelled data for training!

*** Observations:**

1. Activation energy with for a threshold barrier
2. Sensitive certain electric stimulus
3. Activation pattern changes from repeated inputs

Network details

- **Loss function:** $$L_{vp}(\{s\}) = -\sum_{\{s\}}\sum_{v\in\mathcal{V}} e^{-d(v,v_s)/\sigma}\log P_v(s;c_s)$$

- **Network structure (Based on AlexNet):**

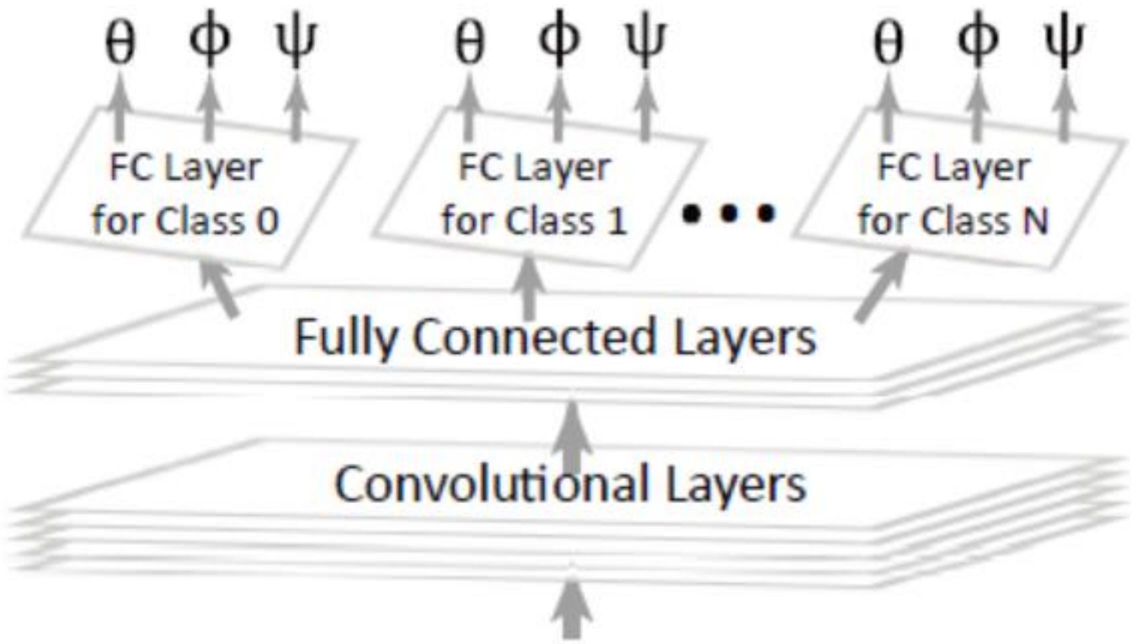

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

More experimental results (Supp.)

- **Azimuth Viewpoint Estimation**

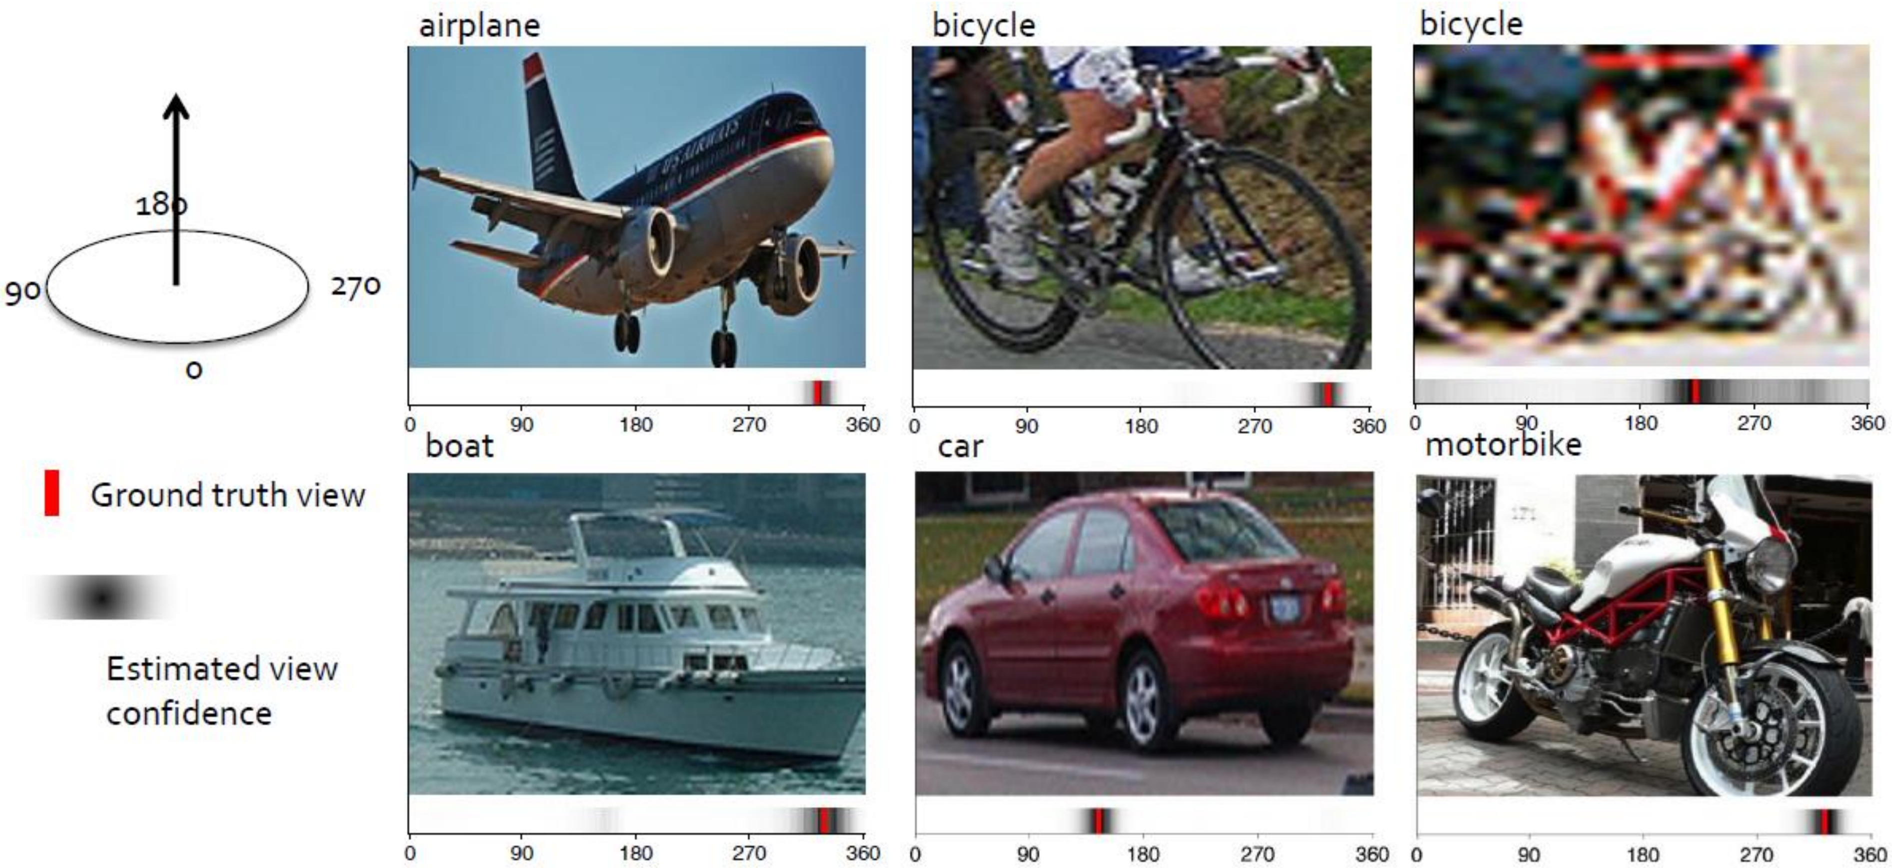

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views
Author's Slides from http://ai.stanford.edu/~haosu/slides/iccv_renderforcnn_website.pdf (ICCV15 Oral)

Quiz

- **Which sampling method will have the lowest variance of the Monte Carlo integration?**

  1. Random Sampling
  2. Gaussian-jittered Sampling $\sigma = 1.0$
  3. Gaussian-jittered Sampling $\sigma = 0.1$
  4. Halton sequence

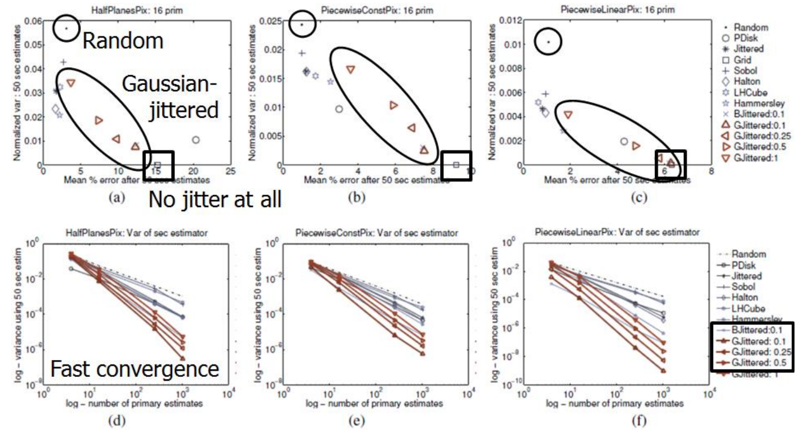

- **Which representation can have the largest number of free parameters to learn from data?**

  1. Gaussian Mixture Model (GMM)
  2. Support Vector Machine (SVM)
  3. Convolutional Neural Network (CNN)
  4. Adaptive Boosting (AdaBoost)

High-cost Label Acquisition **High-capacity Model**

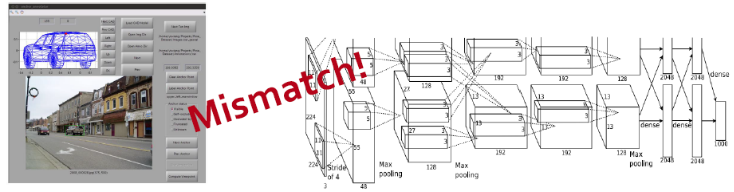

30K images with viewpoint labels in PASCAL3D+ dataset [Xiang et al.]

60M parameters. AlexNet [Krizhevsky et al.]

How to get MORE images with ACCURATE viewpoint labels?